ARENBERG DOCTORAL SCHOOL
Faculty of Engineering Science

SCIENCES AND TECHNOLOGY SECTOR
Ecole Polytechnique de Louvain

Exact Performance Analysis of Fundamental First-Order Optimization Methods

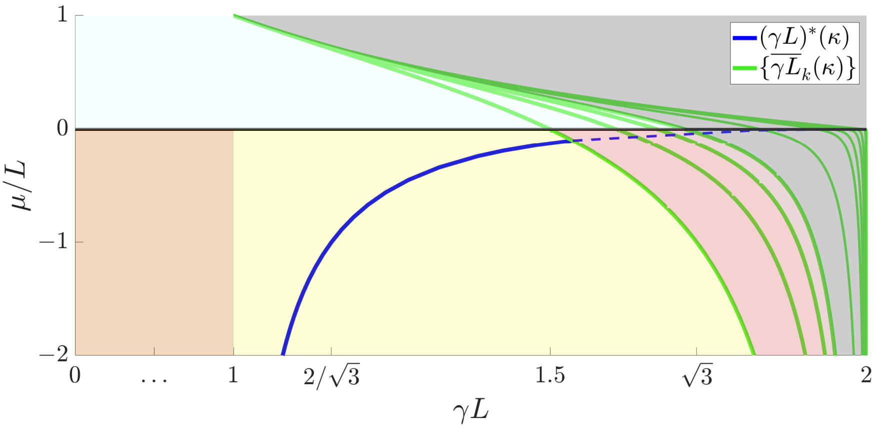

Teodor Rotaru

Dissertation presented in partial fulfillment of the requirements for the degree of Doctor of Engineering Science (PhD): Electrical Engineering

These présenté en vue de l'obtention du grade de docteur en sciences de l'ingenieur et technologie

Supervisors:
Prof. dr. Panagiotis Patrinos (KU Leuven)
Prof. dr. François Glineur (UCLouvain)

February 2026

EXACT PERFORMANCE ANALYSIS OF FUNDAMENTAL FIRST-ORDER OPTIMIZATION METHODS

Teodor ROTARU

Supervisors:
Prof. dr. Panagiotis Patrinos
(KU Leuven)
Prof. dr. François Glineur
(UCLouvain)

Members of the
Examination Committee:
Prof. dr. Adhemar Bultheel, chair
(KU Leuven)
Prof. dr. Nick Vannieuwenhoven
(KU Leuven)
Prof. dr. Aritra Konar
(KU Leuven)
Prof. dr. Julien Hendrickx
(UCLouvain)
Prof. dr. Nelly Pustelnik
(ENS Lyon)
dr. Adrien Taylor
(INRIA Paris)

Dissertation presented in partial fulfillment of the requirements for the degree of Doctor of Engineering Science (PhD): Electrical Engineering

These présenté en vue de l'obtention du grade de docteur en sciences de l'ingenieur et technologie

February 2026

© 2026 Teodor Rotaru
Uitgegeven in eigen beheer, Teodor Rotaru, Kasteelpark Arenberg 10, bus 2446, B-3001 Leuven (Belgium)

Alle rechten voorbehouden. Niets uit deze uitgave mag worden vermenigvuldigd en/of openbaar gemaakt worden door middel van druk, fotokopie, microfilm, elektronisch of op welke andere wijze ook zonder voorafgaande schriftelijke toestemming van de uitgever.

All rights reserved. No part of the publication may be reproduced in any form by print, photoprint, microfilm, electronic or any other means without written permission from the publisher.

To God.

To Ius and Nadi.

To Dada, Si and Ana.

Preface

I began this work intrigued by what a modern tool – PEP – seems to promise as a kind of *Holy Grail* in optimization: automated convergence rates for optimization methods. That is what I tried to pursue over the years. However, I underestimated the path from numbers and nice plots to actual formulas. I still remember that, when applying for a doctoral position, I wrote in my motivation letter that I like writing proofs. This was true (and still is, depending on the day), but I did not realize how many proofs I would end up writing. As my father used to tell me when I was a kid: even if you love pizza, you might start to dislike it after eating it for weeks.

My OCD and perfectionism kept me chasing patterns and proofs for months (and sometimes years) in each chapter, even when I was told it might be better to move on. Without them, this thesis might be more scattered and more practically oriented, with more numerical experiments and a few more analyses that are, in hindsight, less tight than I wanted. Finishing my PhD feels like reaching the end of a journey, sometimes a race against the clock, against low energy, and against growing responsibilities. But other races are still to come. For now, I feel it is my responsibility to thank everyone who helped build the person writing these lines: Dr. me.

Supervisors. When I started my PhD, I saw a meme saying that the supervisor matters more than the research topic or even the student. I thought it was a silly joke and could not possibly be true. After a few years, I understood the importance of such a great academic parent and how much of a guide they can be in the highly challenging and demanding world of research. How lucky I was to have not one, but two such guides: Prof. Panos Patrinos and Prof. François Glineur!

When I applied to the open position with Panos, I was merely testing the waters: my background did not fully match the requirements, beyond a strong

interest in numerical methods. Still, he believed in me from the beginning, waited almost a year until I could start, and then supported me with his broad knowledge, spot-on references, and clear guidance on how to fill my gaps. I also deeply appreciated the friendly and smart-working oriented atmosphere in his group and the rare balance between freedom in research and thoughtful supervision.

I cannot help admiring how François always seemed to be several steps ahead of my research, with intuitions that felt obvious only later, and that I would reach myself weeks later, often after other additional meetings. Many times I challenged his ideas, chose my own (often stubborn) path, and got stuck – until I had to acknowledge that he was right *again*. Beyond that, I deeply appreciated his genuine care: making me feel welcome, solving countless administrative issues, and introducing me to the fascinating world of PEP, where he has been one of the main contributors. Even with my strong preference for truly last-minute submissions, he was always there, offering invaluable help. I will always remember our long whiteboard meetings, playing with equations and brainstorming about worst-case behaviours – moments that made my childhood dream of research feel real.

A big and sincere thank you to both of you!

Committee. I would like to thank all members of my examination committee for their careful reading of the manuscript and for the valuable feedback that helped improve this work. I am grateful to Prof. Nick Vannieuwenhoven for his out-of-the-box remarks during the yearly presentations, which often challenged my way of thinking. I thank Prof. Nelly Pustelnik and Prof. Aritra Konar for their constructive comments and suggestions, which directly contributed to improving the quality and clarity of this manuscript. I thank Prof. Julien Hendrickx for creating an atmosphere of trust during our interactions in LLN, which reinforced my confidence and optimism in my research. Finally, I thank Dr. Adrien Taylor for the many insightful discussions we had; as I often told him, each of them gave a noticeable boost to my research. I also thank Prof. Adhemar Bultheel for chairing my defenses.

Institutions and support. I would like to thank the support provided by the Global PhD Partnership KU Leuven–UCLouvain. Whenever I needed help, I received it at both universities. In particular, I thank Elsy, Ida, John, and Aldona (at Stadius), and Marie-Christine and Pascale (at INMA).

Colleagues. There is a high risk that I will miss some people who deserve credit; please accept my apologies in advance!

I must thank Puya for years of (caring) mockery: “drop the numbers and write real proofs” or “remove the epilepsy-generating plots”. While I might agree with the first part, the second is non-negotiable (Giovanna clearly liked them). Miguel, you made our move to Leuven much more pleasant, and I am glad we still keep in touch. Manu, thank you for our silly discussions about optimization, industry, and life, for the funny accents, and for the good times, whether in the office, cleaning out the RV in the middle of nowhere, or laughing so hard at a conference that we nearly ruined the session. Robin K., you inspired me with your way of balancing PhD life and fatherhood. Renzi, we started the PhD almost at the same time, and I learned a lot about perseverance from your example. Awesome Brecht, the weak-Minty engineer: one day we will finish PEPing some nice 72-regimes proof, hopefully in some remote location that only you can find. Mathijs, despite hosting my birthday at your place and listening to my silly jokes for years, you still seem to enjoy my presence; a true leader (thanks Celine & Julia). Peter, the rigorous all-rounder, thank you. Pieter, even your C++ skills could not prevent political overflow; thank you for our evening office talks. Jan, I am disappointed that you managed to go two whole months without a strong paper in a top venue during your first year of PhD; truly impressive. Kosta, even Siberian enemies can become good friends. Alex, thank you for your genuine care for others, even when your brain generates one research idea per minute (and yes, you can always share them). Leander, I truly admire your “nothing can stop me” attitude. Jia, I am still running on your Chinese water. Matisse, you’ll nail it! I am also grateful to Pourya, Jean-Pierre, Robin B., and Francesco.

I would also like to thank the other current and former Stadius members, with whom I played football, beach volleyball, padel, tennis, or badminton; attended concerts and parties; and enjoyed many coffee-break discussions that made days at ESAT genuinely enjoyable. Raphael, your constant enthusiasm for almost everything is contagious. Angeliki, our perfectly balanced complaint-and-relief conversations were always on point. Simon, looks like your planning obsession is actually useful sometimes. Thomas, I finally met someone more competitive than I am. Many thanks also to Amir, Cem, Charles, Charlotte, Guanchun, Helene, Joran, Konstantinos, Lennert, Luis, Maarten, Manos, Miguel, Nefely, Nico, Rob, Robin W., Stijn, Tim, and Zander.

Even though I spent less time at INMA, I always felt welcome whenever I was there. Many thanks to my office mates Philémon, Hazan and Florentin. I also felt right at home among the other PEPers, and I was lucky to discuss my research with Nizar (and my best wishes to Appoline as well), Anne, Sébastien, and Yassine, as well as with the fresh generation: Pierre, Erwan, and Martina. I also enjoyed exchanging with non-PEP optimizers, in particular Anton, Dânâ, Guillaume, Mihai, Nikita, and Sofiane. Finally, thank you to the young and highly motivated professors Geovani and Estelle for the many pleasant discussions.

Special thanks to Moslem and Hadi: you were both path-openers and the fuel that pushed me to publish my work faster; I am very grateful to have met you both.

Teachers. During my school years, I was lucky to meet professors who had a dramatic (in the best possible sense) impact on my life: Daniela Davidescu (who boosted my love for maths and travelling), Monica Cristea (who taught me that poetry is the grammar for life, like maths is to science), Dorel Haralamb (who shaped my critical eye and deepened my understanding of physics and human nature).

I would also like to thank my professors from University Politehnica of Bucharest, who pushed me to study more, and (almost inevitably) nudged me towards a PhD: Bogdan Ciubotaru (deep thanks for continuous friendship and guidance), Bogdan Dumitrescu (the professor who got me hooked on numerical methods), Ion Necoară (my first contact with optimization), Cristian Oară (a leader who insists on mathematical rigour in engineering), and Florin Stoican (for work ethic and dedication to students).

I must thank the other *Duamne* in the group of ACSE fans: Andrei, Denis, and Sabin — my doctoral brothers — who were role models in different aspects of research, teaching, and dedication to improving academia.

I also want to thank Ionela Prodan, who helped me enormously during the transition from my Master's to my PhD—and I apologise for not choosing the applications-and-control path. Ionuț Farcaș was a tremendous help during some of the hardest periods of my academic life, especially during COVID, always believing in me and supporting me despite the strong evidence that I was close to failing.

Friends. Elvira, I am happy that, after so many years of questioning together, you now know what you want to do in life. I promise not to betray you: I will keep searching. Delia and Bogdan, you were the hidden treasure we found after moving to Leuven; thank you. Mister Alex, you kept me on track by constantly asking how I was doing, and I am also forever in debt to you for training me to work late at night. Bogdan, I am very happy that our friendship is now 20 years old, built on chess, maths, and coding; thank you for all your help over the years. Vlad, you are the maestro showing that true passion can break the

established PhD routes, because chess *rullz*. Andrei and Ana, wait for us. I also want to thank Anca, Bea, Cici, Corina, Dan, Denisse, and Maria.

In memoriam. I want to remember those I miss: Bunelu (find joy while staying serious), Mau (strong life lessons from when I was a kid), Mamaia (a touch of severity, and a lot of love), Olguța, and Vasile Corduneanu (chess-based lessons about how the world works).

Family. However, the first professor I ever interacted with is my father. A day-to-day physics teacher, he is the all-rounder who made me feel bored in school, the person I always felt had the answers, and the one who opened my eyes to the beauty of maths even as a pre-schooler. I know I have disappointed you by not following the predicted *tomberonist* career, or your second proposal, chasing goats in the hills. To be fair, that might indeed have been a less stressful learning path.

To my mother, *mersîc* is not enough. Always there, always present, day and night: for every silly request or joke I made, you still replied with love and help. People hide many things underground, but those bricks of trust are the strongest foundation one can have in life. Țucam pupam!

Oh, SisMătușAna, I am happy you finally stopped chasing me through life in order to prove you can do better. This time you were also smarter. And honestly, you always have been – even back when you let me steal your bank, when we played with knives, or when we were indoor, full-pipeline, firefighters (yes, we survived). Also, thanks to Cristi!

I also thank my special aunts, Mona and Lila, for their constant support, encouragement, massages, and their always-warm thoughts. I also want to thank my parents-in-law, Emi and Dana Andronic, and my sister-in-law Andra for their support.

The last years of my PhD, although among the hardest, were also full of joy thanks to Nadia's smile. Even when we did not see each other for days (the intersection of our wake-up schedules was empty), you always reminded me what matters in life—not some edge-case detail in an intricate proof, but love, in full.

Lastly, in the traditional spirit of *blaming* the wife, I(i)ustina is the guiltiest person *responsible* for this PhD: you found the position, (almost) submitted the application for me, insisted on moving from Munich to Leuven, and we even ended up working on the same floor for our PhDs. You supported me even when I showed more empathy for PEP simulations than I did for you, in my

never-ending search for patterns. You are a great wife, mother, and researcher, and I am very proud of you. We went through the rough PhD years **together**. Mulțuteiu mult, (soon Dr.) Ius!

Teodor Rotaru
Leuven, February 2026

Popularized Abstract

Computers constantly solve optimization problems, which means finding the best possible solution according to some (well-defined) criterion. Many tasks use first-order methods, which follow the “gradient”, or the steepest downhill direction, to find the answer. The most famous example, and the base for many of these methods, is gradient descent, which is a key focus of this thesis. This method is like a hiker always walking downhill to find a valley’s lowest point.

Standard optimization research often focuses on smooth and convex problems, where gradients change gradually rather than abruptly and the surface is bowl-shaped and with a single minimum. This work goes beyond that setting to study smooth weakly convex problems, where curvature can occasionally be negative (sine-like), leading to multiple minima yet still allowing reliable progress with gradient-based methods. Such problems commonly arise in machine learning, where complex models produce structured but nonconvex landscapes.

While we know these methods work, we often only have a general idea of their speed. We might know the distance to the optimal solution shrinks roughly like inverse proportionally to the number of iterations, but this hides the constants determining the precise performance; for example, an algorithm can be guaranteed to be twice as fast. This thesis finds the exact speed limits that cannot be improved for gradient descent and other key algorithms. This shows their true, guaranteed theoretical performance, which may be reached in a worst-case scenario.

To find these limits, the research uses methods inspired by a modern technique called performance estimation problem (PEP). This approach uses a computer to figure out the absolute worst case scenario for an algorithm. Understanding this worst-case limit is vital because it marks the boundary of what these methods can ever achieve. This deepens our knowledge of the mathematical engines that drive modern technology, from artificial intelligence to engineering.

Gepopulariseerde Samenvatting

Computers lossen voortdurend optimalisatieproblemen op: ze zoeken de best mogelijke oplossing volgens een duidelijk criterium. Veel taken gebruiken eerst-orde methoden, die de gradiënt volgen - de steilste afwaartse richting - om het optimum te vinden. Het bekendste voorbeeld, en basis voor veel van deze methoden, is gradient descent, het hoofdonderwerp van deze thesis. De methode lijkt op een wandelaar die steeds bergafwaarts gaat om het laagste punt van een vallei te bereiken.

Traditioneel richt onderzoek zich op gladde, convexe problemen, waarbij gradiënten geleidelijk veranderen en het oppervlak komvormig is met één minimum. Deze thesis bestudeert gladde zwak-convexe problemen, waar de kromming soms negatief kan zijn (zoals bij een sinus). Zulke oppervlakken hebben meerdere minima, maar laten nog steeds betrouwbare voortgang toe met gradiënt-gebaseerde methoden. Dergelijke problemen komen veel voor in machine learning, waar complexe modellen niet-convexe landschappen creëren.

We weten dat deze methoden werken, maar hun snelheid is vaak slechts globaal bekend. Meestal daalt de afstand tot het optimum ongeveer omgekeerd evenredig met het aantal iteraties, maar de precieze constanten blijven verborgen. Zo kan een algoritme theoretisch twee keer zo snel zijn. Deze thesis bepaalt de exacte snelheidsgrenzen voor gradient descent en verwante algoritmen - grenzen die niet verder te verbeteren zijn. Dit onthult hun ware, gegarandeerde prestaties.

Om deze limieten te vinden, gebruikt het onderzoek technieken geïnspireerd op het performance estimation problem (PEP). Deze methode laat een computer het slechtst mogelijke scenario voor een algoritme berekenen. Inzicht in deze worst-case limiet is essentieel: het markeert wat dergelijke methoden maximaal kunnen bereiken. Zo verdiepen we ons begrip van de wiskundige motoren achter moderne technologie - van kunstmatige intelligentie tot ingenieurswetenschappen.

Abstract

In this thesis are derived exact worst-case convergence bounds for three canonical first-order optimization methods: gradient descent, proximal gradient descent (or difference-of-convex algorithm), and Douglas–Rachford splitting, applied to smooth, weakly convex functions.

While the asymptotic behaviour of these algorithms is well established, traditional analyses often rely on big-O notation, which hides the constants that determine their true guaranteed convergence speed. This work goes beyond asymptotic descriptions to establish tight, non-improvable bounds that characterize the exact theoretical performance limits of these methods. Such bounds provide a clearer understanding of the capabilities and limitations of first-order methods in deterministic smooth optimization.

The analysis is fully analytical and interpretable, inspired by the Performance Estimation Problem (PEP) framework. The proofs stand on their own and do not rely on computer-aided verification. Our goal is to provide transparent derivations that reveal the underlying structure of both the algorithms and the analysis, typically in weakly convex settings.

The main results include a complete worst-case analysis of gradient descent with constant stepsizes. This analysis captures its behaviour across weakly convex, convex, and strongly convex regimes, showing how the rates continuously interpolate between different curvature assumptions and providing optimized stepsize schedules, particularly those independent of a priori fixed iteration counts. Extending to constrained minimization, proximal gradient descent is examined through its equivalence with the difference-of-convex algorithm, a simpler parameter-free approach. The analysis provides tight performance bounds on the convergence to critical points. It also extends previous works by introducing the novelty of accommodating weakly convex functions in the second, subtracted term. This yields the broadest admissible range of proximal gradient descent stepsizes and reveals a continuous bridge between smooth nonconvex and convex regimes, with tight rates in all cases. Finally, new exact bounds are established for the Douglas–Rachford splitting method applied to smooth, weakly convex plus prox-bounded functions, covering the entire range of stepsizes and relaxation parameters that ensure convergence.

Together, these results offer a unified and rigorous characterization of the worst-case performance of fundamental first-order methods. Although such worst cases rarely occur in practice, understanding them clarifies the theoretical boundaries of algorithmic behaviour and strengthens the connection between optimization theory and its practical applications.

Beknopte samenvatting

In dit proefschrift worden exacte worst-case-convergentiegrenzen afgeleid voor drie canonieke optimalisatiemethoden van de eerste orde: gradient descent, proximale gradient descent (of het difference-of-convex-algoritme) en Douglas–Rachford-splitting, toegepast op gladde, zwak convexe functies.

Hoewel het asymptotische gedrag van deze algoritmen goed is vastgesteld, steunen traditionele analyses vaak op big-O-notatie, die de constanten verbergt die hun werkelijk gegarandeerde convergentiesnelheid bepalen. Dit werk gaat verder dan asymptotische beschrijvingen en stelt strakke, niet te verbeteren grenzen vast die de exacte theoretische prestatielimieten van deze methoden karakteriseren. Dergelijke grenzen bieden een duidelijker inzicht in de mogelijkheden en beperkingen van eerste-ordemethoden in deterministische gladde optimalisatie.

De analyse is volledig analytisch en interpreteerbaar, geïnspireerd door het Performance Estimation Problem (PEP)-raamwerk. De bewijzen staan op zichzelf en steunen niet op computerondersteunde verificatie. Ons doel is transparante afleidingen te geven die de onderliggende structuur blootleggen van zowel de algoritmen als de analyse, doorgaans in zwak convexe omgevingen.

De belangrijkste resultaten omvatten een volledige worst-caseanalyse van gradient descent met constante stapgroottes. Deze analyse beschrijft het gedrag in zwak convexe, convexe en sterk convexe regimes, toont hoe de snelheden continu interpoleren tussen verschillende krommingsaannames en levert geoptimaliseerde stapgrootteschema's, in het bijzonder schema's die niet afhankelijk zijn van een vooraf vastgelegd aantal iteraties. Uitbreidend naar gecontrainde minimalisatie wordt proximale gradient descent onderzocht via zijn equivalentie met het difference-of-convex-algoritme, een eenvoudigere parameterloze aanpak. De analyse geeft strakke prestatiegrenzen voor de convergentie naar kritieke punten. Zij breidt ook eerder werk uit door de nieuwigheid te introduceren dat zwak convexe functies ook in de tweede,

afgetrokken term kunnen worden toegelaten. Dit levert het breedst toelaatbare bereik van stapgroottes voor proximale gradient descent op en onthult een continue brug tussen gladde niet-convexe en convexe regimes, met strakke snelheden in alle gevallen. Tot slot worden nieuwe exacte grenzen vastgesteld voor de Douglas–Rachford-splittingmethode toegepast op gladde functies van het type "zwak convex plus prox-begrensd", waarbij het volledige bereik van stapgroottes en relaxatieparameters dat convergentie garandeert wordt afgedekt.

Samen bieden deze resultaten een verenigde en rigoureuze karakterisering van de worst-caseprestaties van fundamentele eerste-ordemethoden. Hoewel zulke worst cases zich in de praktijk zelden voordoen, verheldert hun begrip de theoretische grenzen van algoritmisch gedrag en versterkt het de verbinding tussen optimalisatietheorie en haar praktische toepassingen.

List of Abbreviations

ADMM Alternating Direction Method of Multipliers. 11

DC Difference-of-Convex. 104

DCA Difference-of-Convex Algorithm. 10

DRE Douglas–Rachford Envelope. 174

DRS Douglas–Rachford Splitting. 11

GD Gradient Descent. 9

l.h.s. left-hand side. 13

l.s.c. lower semicontinuous. 14

PEP Performance Estimation Problem. 3

PGD Proximal Gradient Descent. 10

PPA Proximal Point Algorithm. 145

PSD Positive Semidefinite. 13

r.h.s. right-hand side. 13

s.t. subject to. 209

SDP Semidefinite Program. 17

List of Symbols

$\mathbb{R}$	Real numbers
$\bar{\mathbb{R}}$	Extended real numbers $\mathcal{R} \cup \{\infty\}$
$\mathbb{R}^{n\times m}$	Space of real $n \times m$ matrices
$\emptyset$	Empty set
$(\cdot)$	Placeholder for the argument of a function
$[x]_+$	Positive part of x: $\max\{0, x\}$
$[x]_-$	Negative part of x: $\min\{0, -x\}$
$A^\top$	Transpose of matrix A
$A_{(i,j)}$	The (i, j)-entry in matrix A
$A_{(i_0:i_1, j_0:j_1)}$	The submatrix of A with rows i_0 to i_1 and columns j_0 to j_1
$\lambda_{\min / \max}(A)$	The minimum / maximum eigenvalue of matrix A
I	Identity matrix
O	Zero matrix
δ_X	Indicator function of set X
$f\colon A \to B$	Single-valued function
$\operatorname{dom} f$	Domain of extended-real-valued function f
$\operatorname{range} f$	Range of extended-real-valued function f
$\operatorname{ri} C$	Relative interior of set C
x_* or x^*	Local / global optimal solution

∇f	Gradient of f
$\nabla^2 f$	Hessian of f
id	The identity mapping
$T\colon A \rightrightarrows B$	A set-valued operator T
T^{-1}	Inverse of operator T
prox_f	Proximal mapping of f
$\mathrm{dist}(x, X)$	Distance of x from set X (in Euclidean norm)
∂f	Convex (Fenchel) subdifferential of f
$\hat{\partial} f$	Regular subdifferential of f
$\partial_L f$	Limiting (Mordukhovich) subdifferential of f
$\mathcal{O}(\cdot)$	Big-O Bachmann-Landau notation
$o(\cdot)$	Little-o Bachmann-Landau notation
$\mathcal{F}$	Function class
$\mathcal{F}_{\mu,L}$	Class of smooth functions whose curvature lies in $[\mu, L]$
$\mathcal{F}_{\mu,\infty}$	Class of nonsmooth functions with lower curvature bounded by μ
κ	Curvature's ratio $\frac{\mu}{L}$ ("inverse condition number")
$PB^{w}_{x,y}$	Interpolation conditions for prox-bounded functions
$\mathcal{I}$	Index set
$\mathcal{T}$	First-order oracle / triplets
g_i	(Sub)Gradient of f evaluated at iterate x_i
f_i	Function f evaluated at iterate x_i
i or k	Iteration count
N	Number of iterations
γ	Stepsize
γL	Normalized stepsize
$(\gamma L)^*$	Optimized normalized stepsize
$\overline{\gamma L}$	Stepsize thresholds

Contents

List of Figures

List of Tables

List of Statements

Chapter 1

Introduction

Numerical optimization is a field at the intersection of applied mathematics and computer science that studies how to find the best possible solution to a problem using numerical methods. It aims to minimize or maximize an objective function, usually subject to certain constraints. In many cases, exact analytical solutions to such problems are difficult or even impossible to obtain. Therefore, they are solved numerically using optimization algorithms. Many real-world problems can be expressed as optimization problems, for example, training machine learning models or designing engineering systems.

First-order methods are optimization algorithms that update the solution using only first-order information about the objective function, namely gradient or subgradient evaluations and proximal operators of the functions that make up the objective.

For smooth optimization problems, gradient-based methods decrease the function value at each step by following the direction of the negative gradient. The most fundamental of these is gradient descent, which forms the basis of many other first-order algorithms. Proximal, accelerated, and splitting methods can all be seen as generalizations or extensions of gradient descent that handle additional problem structures, such as nonsmooth terms, constraints, or separable objectives. These methods are popular because they are simple, require little memory, and scale well to large problems. They are especially useful for large-scale problems where second-order methods, which rely on Hessian information, are too expensive.

Proximal splitting methods are first-order algorithms designed for structured problems where the objective is written as the sum of smooth and possibly nonsmooth components. The smooth part is handled with gradient steps, while the nonsmooth part, often representing constraints through indicator functions, is handled with its proximal operator. The two updates are then combined through fixed-point iterations.

This work focuses on non-accelerated first-order methods for nonlinear and deterministic objective functions that can be written as composite problems with at least one smooth term. In machine learning, training models such as logistic regression, support vector machines, and neural networks often involves minimizing smooth loss functions using first-order methods. In engineering, smooth optimization appears in control design, signal processing, and parameter estimation, where gradients can be computed efficiently. These applications require scalable algorithms that maintain predictable convergence behaviour, making first-order methods a natural and practical choice.

In this thesis, while sometimes also considering convex problems, we focus on smooth weakly convex (or hypoconvex) functions, a broad and structured class of nonconvex objectives that satisfy a controlled curvature condition keeping them close to convex behaviour. In such cases, first-order methods generally guarantee convergence to stationary points, which may correspond to local minimizers.

The efficiency of optimization methods is evaluated through their iteration complexity, which measures how fast an algorithm approaches the optimal solution. This progress is typically quantified by the decrease in the objective value or the gradient norm. Complexity bounds describe how many iterations or gradient evaluations are required to achieve a desired accuracy, providing convergence rates, for example, of order $\mathcal{O}(1/N)$, where N is the number of iterations the algorithm performs. These bounds are essential because they offer a theoretical measure of algorithmic efficiency, allowing one to compare methods, predict their behaviour, and design improved algorithms. In this way, complexity analysis connects the mathematical theory of optimization with its practical performance in applications.

A *tight* (or *exact*) convergence rate of an algorithm is a performance bound that is actually attained by some function within the considered class. Obtaining tight convergence rates with explicit constants is important for a precise understanding of algorithmic performance. While asymptotic rates such as $\mathcal{O}(1/N)$ describe the general speed of convergence, they often omit the constants that determine the actual rate of progress. Two algorithms sharing the same asymptotic rate may differ substantially in practice because of these constants, and even comparisons between methods with different asymptotic orders can change once constants are accounted for. Tight bounds with explicit constants therefore provide a more accurate and fair basis for comparison, helping researchers assess efficiency, identify optimal methods, and understand fundamental performance limits. Such results also strengthen the connection between theoretical analysis and practical observations, offering a clearer picture of how first-order methods perform in real applications and in challenging scenarios equal to or close to worst cases.

There is a gap between well-established convergence proofs and exact ones (with explicit constants) for which one can provide worst-case problem instances. The seminal work [44] initiated a systematic, computer-aided approach to address this gap, called the performance estimation problem (PEP). The methodology in [44] has been extended to generate numerical certificates for the tightest possible bounds, together with worst-case instances and a principled way to combine inequalities to generate a mathematical proof [116, 113, 115, 56, 124]. With this tool, it is easy to verify whether a performance bound from the literature is improvable for specific parameter values such as the number of iterations, curvature parameters describing the problem class, or tuning hyperparameters like stepsizes. However, the main difficulty lies in obtaining analytical worst-case values and proofs, due to intricate dependencies between all these parameters.

Consequently, the main research goal of this thesis is formulated as follows:

Obtain exact, non-improvable, convergence bounds for canonical gradient- and proximal-based methods, with a focus on weakly convex optimization.

Our derivations are inspired by the PEP framework; however, several of our arguments differ from the computer-generated proofs, which are often difficult to interpret. Besides providing non-improvable bounds, we aim to present clear and principled proofs, typically in the form of descent-lemma-type arguments. This approach helps us better understand special cases, reveal patterns, and identify connections between different optimization problems, leading to more unified analyses.

In this work, we consider gradient descent (GD), proximal gradient descent (PGD), and Douglas–Rachford splitting (DRS) as canonical examples of first-order algorithms in smooth nonconvex settings. These methods rely solely on first-order information, such as gradients or proximal mappings, and lie at the foundation of modern optimization theory. Gradient descent serves as the basic approach for smooth problems, while proximal gradient descent extends it to composite problems that include a smooth term and a simple nonsmooth regularizer. We analyse the performance of PGD by using an extended version of the difference-of-convex algorithm (DCA). The Douglas–Rachford splitting method further generalizes this framework to problems involving the sum of multiple structured functions. Together, these algorithms represent fundamental building blocks for understanding the behaviour and convergence properties of first-order methods in nonconvex settings.

1.1 Motivating example for GD, PGD/DCA, DRS

To motivate the theoretical developments of this thesis, we introduce a running example that progressively incorporates the structures targeted by gradient descent (GD), proximal gradient descent (PGD) (and its DCA viewpoint), and Douglas–Rachford splitting (DRS). The goal is not to emphasize a particular application, but to provide a concrete template that clarifies (i) why each algorithm is natural for a given objective structure and (ii) why different stationarity measures arise in each chapter.

Indexing conventions. To remain consistent with the subsequent chapters, we use subscripts for the GD and DRS iterates (e.g., x_k and s_k), and superscripts for the PGD/DCA iterates (e.g., x^k). This distinction is purely notational.

Step 1. A smooth (possibly weakly convex) objective f and GD (Chapter 3). We begin with the smooth unconstrained problem

$$\underset{x\in\mathbb{R}^d}{\text{minimize}}\, f(x),$$

where f is differentiable with Lipschitz continuous gradient and, in the nonconvex regime of interest, may further belong to the bounded-curvature class $\mathcal{F}_{\mu,L}$ (with curvature bounds μ and L). A prototypical example is the quadratic model

$$f(x) := \frac{1}{2}\|Ax-b\|^2 - \frac{\rho}{2}\|x\|^2, \quad A\in\mathbb{R}^{m\times d},\ b\in\mathbb{R}^m,\ \rho\geq 0,$$

for which $\nabla f(x) = A^\top(Ax-b)-\rho x$ and $\nabla^2 f(x) = A^\top A - \rho I$. Hence $f\in\mathcal{F}_{\lambda_{\min}(A^\top A)-\rho,\ \lambda_{\max}(A^\top A)-\rho}$, so f is smooth and may be weakly convex when $\lambda_{\min}(A^\top A)-\rho<0$. In this setting, GD iterates as

$$x_{k+1} = x_k - \gamma\nabla f(x_k),$$

and the natural stationarity measure is the gradient norm $\|\nabla f(x)\|$, which motivates the performance criteria in Chapter 3:

$$\min_{0\leq k\leq N}\{\|\nabla f(x_k)\|^2\}.$$

In the convex setting (for a standard admissible range), the gradient norm is nonincreasing, so the best one is attained at the last iterate. As initial condition, we consider either the optimality gap $f(x_0) - f_*$ or the last iterate gap $f(x_0) - f(x_N)$.

Step 2. A composite objective $F = \varphi + h$ and stationarity measures for PGD/DCA (Chapter 4). To model nonsmooth regularization (e.g., sparsity), we consider the composite problem

$$\underset{x \in \mathbb{R}^d}{\text{minimize}}\, F(x) \coloneqq \varphi(x) + h(x),$$

where $\varphi \in \mathcal{F}_{\mu_\varphi, L_\varphi}$ is differentiable with Lipschitz gradient (possibly weakly convex) and $h \in \mathcal{F}_{\mu_h, L_h}$ is proper, lower semicontinuous, and proximable (e.g., $h(x) = \lambda \|x\|_1$; when $\mu_h = 0$ and $L_h = \infty$, h is convex). Throughout the PGD/DCA analysis we assume a feasible initialization, namely $x^0 \in \operatorname{dom} h$ (equivalently $F(x^0) < \infty$). In the running example one may take $\varphi \equiv f$ and add h as a regularizer. Stationarity is expressed by

$$0 \in \nabla \varphi(x) + \partial h(x).$$

With a stepsize $\gamma > 0$, PGD applies

$$x^{k+1} = \operatorname{prox}_{\gamma h}\big(x^k - \gamma \nabla \varphi(x^k)\big)\,, \tag{1.1}$$

We consider two closely related certificates of distance to stationarity. As an initial condition, we assume that the last iterate gap $F(x^0) - F(x^N)$ is bounded. This quantity is always upper bounded by the optimality gap $F(x^0) - \inf F(x)$.

i) **Gradient mapping.** Define the (proximal-)gradient mapping

$$\frac{1}{\gamma}(x^k - x^{k+1}).$$

By the proximal optimality condition, there exists $g_h^{k+1} \in \partial h(x^{k+1})$ such that

$$\frac{1}{\gamma}(x^k - x^{k+1}) = \nabla \varphi(x^k) + g_h^{k+1}.$$

Thus $\|\nabla \varphi(x^k) + g_h^{k+1}\|$ is computable from the PGD step and vanishes at stationary points (under standard qualification conditions). We track

$$\min_{0 \le k \le N} \big\{\|\nabla \varphi(x^k) + g_h^{k+1}\|^2\big\}.$$

ii) **Residual subgradient.** We also measure distance to stationarity via

$$\min_{0\le k\le N}\left\{\|\nabla\varphi(x^k)+g_h^k\|^2\right\},\quad g_h^k\in\partial h(x^k),$$

interpreted for an algorithmically consistent selection of g_h^k (since $\partial h(x^k)$ is set-valued). A canonical choice is induced by the PGD step: for $k\ge 0$ the optimality condition of (1.1) yields

$$g_h^{k+1}=\frac{1}{\gamma}(x^k-x^{k+1})-\nabla\varphi(x^k)\in\partial h(x^{k+1}),$$

This viewpoint is particularly convenient in the equivalent DCA formulation, where subgradient selections are intrinsic to the method.

DCA viewpoint and equivalence. In the DCA analysis, we consider a DC decomposition

$$F(x)=f_1(x)-f_2(x),$$

with associated subgradients $g_1^k\in\partial f_1(x^k)$ and $g_2^k\in\partial f_2(x^k)$. The DCA step

$$x^{k+1}\in\operatorname*{arg\,min}_{w\in\mathbb{R}^d}\{f_1(w)-\langle g_2^k,w\rangle\}$$

is characterized by the optimality condition $g_1^{k+1}=g_2^k$. When at least one component is smooth, we measure near-criticality through the residual gradient $\|g_1^k-g_2^k\|$ and through the iterate gap $\|x^k-x^{k+1}\|$ (equivalently, the gradient mapping up to scaling). To connect PGD and DCA, set

$$f_1(x)\coloneqq h(x)+\frac{1}{2\gamma}\|x\|^2,\quad f_2(x)\coloneqq\frac{1}{2\gamma}\|x\|^2-\varphi(x),$$

so that $F=f_1-f_2$ and the DCA update coincides with the PGD update (1.1).

In the running example with

$$\varphi(x)=\frac{1}{2}\|Ax-b\|^2-\frac{\rho}{2}\|x\|^2,\quad h(x)=\lambda\|x\|_1,$$

we obtain a DC decomposition in which f_1 is (nonsmooth) strongly convex and f_2 is smooth (and possibly weakly convex). Concretely,

$$f_1(x)=\lambda\|x\|_1+\frac{\gamma^{-1}}{2}\|x\|^2,\quad f_2(x)=\frac{\gamma^{-1}+\rho}{2}\|x\|^2-\frac{1}{2}\|Ax-b\|^2.$$

The curvature bounds then translate from the PGD setting

$$\varphi\in\mathcal{F}_{\lambda_{\min}(A^\top A)-\rho,\ \lambda_{\max}(A^\top A)-\rho},\quad h(x)\in\mathcal{F}_{0,\infty}$$

to the corresponding DCA setting

$$f_2\in\mathcal{F}_{\gamma^{-1}+\rho-\lambda_{\max}(A^\top A),\ \gamma^{-1}+\rho-\lambda_{\min}(A^\top A)},\quad f_1\in\mathcal{F}_{\gamma^{-1},\infty}.$$

Step 3: A splitting objective $\varphi = \varphi_1 + \varphi_2$ and DRS (Chapter 5). Finally, to cover problems that split into two proximable terms (possibly including nonconvex but prox-bounded regularizers), we consider

$$\underset{x\in\mathbb{R}^d}{\text{minimize}}\, \varphi(x) := \varphi_1(x) + \varphi_2(x),$$

where φ_1 is smooth (possibly nonconvex) and φ_2 is proper, lower semicontinuous, and prox-bounded (possibly nonconvex). Using the notation of Chapter 5, DRS generates an outer sequence (s_k) and inner iterates (u_k) and (v_k) via

$$\begin{aligned} u_k &\in \operatorname{prox}_{\gamma\varphi_1}(s_k),\\ v_k &\in \operatorname{prox}_{\gamma\varphi_2}(2u_k - s_k),\\ s_{k+1} &= s_k + \lambda(v_k - u_k), \end{aligned}$$

with stepsize $\gamma > 0$ and relaxation parameter λ. Unlike GD and PGD, DRS is not primarily a descent method for φ; its convergence is naturally formulated in terms of fixed-point/progress measures and stationarity residuals associated with

$$0 \in \nabla\varphi_1(x) + \partial\varphi_2(x).$$

Accordingly, we track two residuals:

i) **Iterate progress:**

$$P^{\Delta} := \min_{0\le k\le N} \left\{ \left\|\gamma^{-1}(u_k - v_k)\right\|^2 \right\}.$$

ii) **Residual subgradient:**

$$\begin{aligned} P^{\nabla} &:= \min_{0\le k\le N-1} \{\operatorname{dist}\left(0, \hat{\partial}\varphi(v_k)\right)\}\\ &= \min_{0\le k\le N} \left\{ \left\|\gamma^{-1}(u_k - v_k) - \left[\nabla\varphi_1(u_k) - \nabla\varphi_1(v_k)\right]\right\|^2 \right\}. \end{aligned}$$

As an initial condition, we assume that the Douglas–Rachford envelope gap is bounded, namely $\varphi_\gamma^{\mathrm{DR}}(s_0) - \inf\varphi < \infty$.

Summary. We start from smooth optimization (GD), then study composite objectives $F = \varphi + h$ (PGD/DCA) covering both smooth and nonsmooth regimes, and finally analyze splittings $\varphi = \varphi_1 + \varphi_2$ where φ_1 is smooth and φ_2 is treated by a proximal step (possibly nonconvex but prox-bounded), which motivates DRS. In Table 1.1 we summarize the analyzed problem classes, algorithms and performance measures (including the initial condition).

Table 1.1: Summary of problem classes, methods, performance measures, and initial conditions (bounded quantities). Abbreviations: S=smooth, NS=nonsmooth, w.c.=weakly convex, c.=convex, s.c.=strongly convex.

Ch.	Method	Objective	Assumptions / class	Performance measure	Initial condition
3	GD	$f(x)$	$f \in \mathcal{F}_{\mu,L}$ (S: w.c./c./s.c.)	best gradient $\min_{0\le k\le N} \|\nabla f(x_k)\|^2$	$f(x_0) - f(x_N)$ $f(x_0) - f_\star$
4	PGD	$F(x) = \varphi(x) + h(x)$	$\varphi \in \mathcal{F}_{\mu_\varphi,L_\varphi}$ (S: w.c./c./s.c.); $h \in \mathcal{F}_{\mu_h,L_h}$ proximable (NS)	(i) residual gradient $\min_{0\le k\le N} \|\nabla\varphi(x^k) + \partial h(x^k)\|^2$ (ii) gradient mapping $\min_{0\le k\le N} \|\gamma^{-1}(x^k - x^{k+1})\|^2$	$F(x^0) - F(x^N)$ $F(x^0) - F_{\text{lo}}$
4	DCA	$F(x) = f_1(x) - f_2(x)$	$f_1 \in \mathcal{F}_{\mu_1,L_1}$ (c./s.c.; S/NS) $f_2 \in \mathcal{F}_{\mu_2,L_2}$ (w.c./c./s.c.; S/NS)	(i) residual gradient $\min_{0\le k\le N} \|\partial f_1(x^k) - \partial f_2(x^k)\|^2$ (ii) gradient mapping $\min_{0\le k\le N} \|x^k - x^{k+1}\|^2$	$F(x^0) - F(x^N)$ $F(x^0) - F_{\text{lo}}$
			$f_1,\ f_2$ NS	(iii) Bregman divergence $\min_{0\le k\le N} T(x^k)$	
5	DRS	$\varphi(s) = \varphi_1(s) + \varphi_2(s)$	$\varphi_1 \in \mathcal{F}_{\mu,L}$ (S: w.c./c.) φ_2 l.s.c. prox-bounded	(i) iterate progress $\min_{0\le k\le N} \|\gamma^{-1}(u_k - v_k)\|^2$ (ii) residual subgradient $\min_{0\le k\le N}\{\text{dist}\big(0, \hat\partial\varphi(v_k)\big)\}^2$	$\varphi_\gamma^{\text{DR}}(s_0) - \inf\varphi$

1.2 Thesis contributions and organization

This thesis improves existing theoretical results by providing tight, non-improvable convergence guarantees for all the methods under consideration. Our work can be viewed as a collection of tight-rate results, with analytical formulas that can serve as final exact bounds for inclusion in optimization textbooks. Most of these results are accompanied by complete proofs, while in a few remaining cases they are supported by intuitions and informed guesses confirmed through PEP-based verification.

A tight (exact) performance bound is established in two steps: (i) deriving an *upper bound* on the convergence measure and (ii) constructing a problem instance (a *lower bound*) achieving it. For convenience, the performance measure is often called *(convergence) rate*. This term refers not only to the asymptotic order, such as $\mathcal{O}(\cdot)$, but also to its associated, and often exact, constant.

Chapter 2 introduces the necessary theoretical background, with a focus on explaining the Performance Estimation Problem (PEP) methodology.

The rest of the thesis is organized so that each method is featured in its own dedicated chapter.

Chapter 3: A complete analysis of gradient descent

In the unconstrained setting, we provide a complete study of the canonical optimization method: gradient descent (GD) with constant stepsizes. We derive exact worst-case convergence rates on the best gradient norm of the iterates. Our analysis covers all possible stepsizes and arbitrary upper/lower bounds on the curvature of the objective function, thus including convex, strongly convex and weakly convex (hypoconvex) objective functions.

Among the challenging parts of the analysis, we note the necessity to exploit dependencies between non-consecutive iterates. While this complicates the proofs to some extent, it enables us to achieve an exact full-range analysis of gradient descent for any constant stepsize (covering, in particular, normalized stepsizes greater than one), whereas the literature contained only conjectured rates of this type.

In the nonconvex case, allowing arbitrary bounds on upper and lower curvatures extends existing partial results that are valid only for gradient Lipschitz functions (i.e., where lower and upper bounds on curvature are opposite), leading to improved rates for weakly convex functions.

From our exact worst-case performance bounds, we deduce the optimal constant stepsize for gradient descent for any number of iterations. Leveraging our analysis, we also introduce a new variant of gradient descent based on a unique, fixed sequence of variable stepsizes, demonstrating its superiority in the worst-case over any constant stepsize schedule.

Chapter 4: A comprehensive study of the difference-of-convex algorithm and proximal gradient descent

In constrained settings, the natural extensions of gradient descent are projected and proximal gradient descent (PGD) algorithms, which belong to the framework of proximal splitting methods. They incorporate constraints into a term to which the proximal mapping is applied. Using the equivalence between the PGD iteration and another well-established method, the difference-of-convex algorithm (DCA), we provide an extensive analysis of the latter, which is simpler to perform, and thereby obtain exact bounds for both methods.

DCA is a simple, classical parameter-free algorithm designed to minimize a difference of convex functions. We present a comprehensive convergence analysis of its extension to the case where the second, subtracted function is weakly convex. This generalization allows us to derive performance bounds for larger PGD stepsizes (greater than the inverse of the Lipschitz constant).

Assuming at least one function is smooth, we evaluate the algorithm's performance using the norm of the best (sub)gradient residual or the best (proximal-)gradient mapping. Six distinct regimes are identified, two of them corresponding to the previously studied standard difference-of-convex setting. Our simplified proofs are based on deriving six distinct descent lemmas, inspired by solving the associated performance estimation problems (PEP). These lemmas establish sublinear convergence rates, which are provably tight for three of the regimes in each case. We also conjecture the exact behaviour across any number of iterations in the other regimes, informed by extensive numerical simulations and by connections to the gradient descent method.

Additionally, we propose a new technique to improve the guaranteed worst-case convergence performance through a curvature-shifting strategy, which can transform the subtracted function into a weakly convex one.

Chapter 5: Exact performance bounds for the Douglas–Rachford splitting for smooth plus prox-bounded objective functions

Constraints can be replaced by adding more general nonsmooth terms in the objective function, provided that they satisfy the prox-boundedness condition. In the setting where they are composed with a smooth weakly convex function, we study an alternative proximal algorithm to PGD, namely the Douglas–Rachford splitting (DRS).

Unlike PGD, DRS involves two proximal mapping steps per iteration, which often improves numerical conditioning and empirical performance in nonconvex problems. Moreover, DRS underlies the alternating direction method of multipliers (ADMM), a widely used algorithm in machine learning, signal processing, and control.

We derive explicit performance bounds across the full admissible range of stepsizes and relaxation parameters. Our guarantees quantify either (i) the (best) iterate progress or (ii) the (best) residual subgradient. For both performance measures, we prove three distinct regimes governed by the curvature parameters of the smooth component, all exhibiting sublinear convergence. Numerical evidence from the PEP framework confirms that these bounds are tight for any number of iterations. Finally, we provide practical hyperparameter tuning recommendations derived from the worst-case analysis and validate them on illustrative examples.

Chapter 2

Preliminaries

In this chapter, we introduce the notation and theoretical background that form the basis of the results developed throughout this thesis. Section 2.1 provides the key definitions and mathematical preliminaries used in the analysis. Section 2.2 presents the performance estimation problem (PEP) framework, together with the proof methodology and the notion of interpolation inequalities, which serve as the central analytical tools for deriving exact convergence results.

2.1 Theoretical background

We recall standard definitions and results that are used throughout the thesis (see, e.g., [89, 15, 16, 96, 98]).

The set of real numbers is denoted by $\mathbb{R}$, and the set of extended real numbers by $\bar{\mathbb{R}} \coloneqq (-\infty, +\infty]$. For any $x \in \mathbb{R}$, we define the positive and negative parts as $[x]_+ \coloneqq \max\{x, 0\}$ and $[x]_- \coloneqq \min\{-x, 0\}$, respectively.

The identity operator is denoted by id. A function class is denoted by $\mathcal{F}$. We use l.h.s. and r.h.s. to refer to the left-hand and right-hand sides of an inequality, respectively.

For a matrix A, we denote its entries by $A_{(i,j)}$. A symmetric matrix A is positive semidefinite (PSD), written $A \succeq 0$, if all its eigenvalues are nonnegative. The identity matrix is denoted by I, and its dimension is typically clear from context. For symmetric matrices, the elements below the main diagonal may be indicated by a dot "$\cdot$" when omitted for brevity.

For a function $f\colon \mathbb{R}^d \to \bar{\mathbb{R}}$, we define its *domain* and *range* as

$$\operatorname{dom} f \coloneqq \{x \in \mathbb{R}^d \colon f(x) < +\infty\},$$

$$\operatorname{range} f \coloneqq \{y \in \mathbb{R} \colon \exists x \in \operatorname{dom} f \text{ s.t. } y = f(x)\}.$$

A function is said to be *proper* if it is not identically equal to $+\infty$, and is said to be *lower semicontinuous* (l.s.c.) if $\liminf_{x\to\bar{x}} f(x) \geq f(\bar{x})$ for all $\bar{x} \in \mathbb{R}^d$.

A function f is *convex* if, for all $x, y \in \mathbb{R}^d$ and $\theta \in [0,1]$,

$$f(\theta x + (1-\theta)y) \leq \theta f(x) + (1-\theta)f(y).$$

If f is differentiable, this is equivalent to

$$f(y) \geq f(x) + \langle \nabla f(x), y - x\rangle,$$

for all $x, y \in \mathbb{R}^d$.

Smooth and weakly convex functions

A differentiable function f is called *Lipschitz smooth* (or B-smooth) if its gradient is Lipschitz continuous with constant $B > 0$, i.e.,

$$\|\nabla f(x) - \nabla f(y)\| \leq B\|x - y\|, \quad \forall x, y \in \mathbb{R}^d.$$

Throughout this work, we use a more general characterization of smoothness based on curvature bounds.

Definition 2.1.1 (Bounded curvature class $\mathcal{F}_{\mu,L}$)**.** *Let $L \in \bar{\mathbb{R}}$ and $\mu \in [-\infty, L]$. A function $f\colon \mathbb{R}^d \to \mathbb{R}$ belongs to the class $\mathcal{F}_{\mu,L}$ if it satisfies the following curvature–convexity conditions:*

(i) ***(Upper curvature bound)*** *If $L < +\infty$, then the function*

$$x \mapsto \tfrac{L}{2}\|x\|^2 - f(x)$$

is convex. (If $L = +\infty$, this condition is understood as vacuous, i.e., no upper curvature bound is imposed.)

(ii) ***(Lower curvature bound)*** *If $\mu > -\infty$, then the function*

$$x \mapsto f(x) - \tfrac{\mu}{2}\|x\|^2$$

is convex. (If $\mu = -\infty$, this condition is vacuous, i.e., no lower curvature bound is imposed.)

Intuitively, if $f \in \mathcal{C}^2 \cap \mathcal{F}_{\mu,L}$, then this means that the eigenvalues of its Hessian lie in $[\mu, L]$. Depending on the sign of μ, f is categorized as: (i) weakly convex (or hypoconvex) if $\mu < 0$, (ii) convex if $\mu = 0$, or (iii) strongly convex if $\mu > 0$. Moreover, depending on the sign of L, f is categorized as: (i) concave if $L = 0$, (ii) B-smooth, with $B = \max\{-\mu, L\}$, if $L > 0$, or (iii) strongly concave if $L < 0$.

Weakly convex functions were introduced in [92]; see also [95, 96, 45]. The class of smooth weakly convex functions includes smooth not-necessarily convex functions under the condition $\mu = -L$.

Functions in $\mathcal{F}_{\mu,L}$ are necessarily smooth when $L < \infty$, whereas $\mathcal{F}_{\mu,\infty}$ also includes nonsmooth functions.

For convenience, we define the curvature ratio $\kappa \coloneqq \frac{\mu}{L} \in (-\infty, 1]$, which, when $\mu \geq 0$, can be the inverse condition number.

Subdifferentials and related notions

A set-valued mapping $H \colon \mathbb{R}^n \rightrightarrows \mathbb{R}^m$ is a correspondence from $\mathbb{R}^n$ to the power set $\mathcal{P}(\mathbb{R}^m)$ (the set of all subsets of $\mathbb{R}^m$). Its *graph* is defined as $\operatorname{gph} H \coloneqq \{(x, y) \in \mathbb{R}^n \times \mathbb{R}^m | y \in H(x)\}$.

For $\alpha \in \mathbb{R}$, the α-level set of a function $f \colon \mathbb{R}^d \to \mathbb{R}$ is $\operatorname{lev}_{\leq\alpha} f \coloneqq \{x \in \mathbb{R}^d \colon f(x) \leq \alpha\}$. A function f is *level-bounded* if $\operatorname{lev}_{\leq\alpha} f$ is bounded for all $\alpha \in \mathbb{R}$.

We denote by $\partial f \colon \mathbb{R}^d \rightrightarrows \mathbb{R}^d$ the *convex (Fenchel-Rockafellar) subdifferential* of $f \colon \mathbb{R}^d \to \bar{R}$, where

$$\partial f(x) \coloneqq \{g \in \mathbb{R}^d \,|\, f(y) \geq f(x) + \langle g, y - x\rangle \ \forall y \in \mathbb{R}^d\}.$$

The regular (Fréchet) subdifferential $\hat{\partial} f$ of $f \colon \mathbb{R}^d \to \bar{\mathbb{R}}$ is given by (e.g., see [98])

$$g \in \hat{\partial} f(\bar{x}) \iff \liminf_{\substack{x \to \bar{x} \\ x \neq \bar{x}}} \frac{f(x) - f(\bar{x}) - \langle g, x - \bar{x}\rangle}{\|x - \bar{x}\|} \geq 0.$$

We denote by $\partial_L f \colon \mathbb{R}^d \rightrightarrows \mathbb{R}^d$ the limiting (Mordukhovich) subdifferential of $f \colon \mathbb{R}^d \to \bar{\mathbb{R}}$, which is defined as follows (e.g., see [98, 80]):

$$g \in \partial_L f(\bar{x}) \iff \exists (x_k, g_k)_{k\in\mathbb{N}} \subseteq \operatorname{gph} \hat{\partial} f \text{ such that } (x_k, f(x_k), g_k) \to (\bar{x}, f(\bar{x}), g).$$

According to [98, Theorem 10.1], a necessary condition for local minimality of a point x for a function f is $0 \in \hat{\partial} f(x)$. The domain and range of the subdifferential are

$$\operatorname{dom} \partial f = \{x \in \mathbb{R}^d : \partial f(x) \neq \emptyset\},$$

$$\operatorname{range} \partial f = \cup_{x \in \operatorname{dom} \partial f} \{\partial f(x)\}.$$

The convex conjugate of a l.s.c. function f is

$$f^*(y) \coloneqq \sup_{x \in \operatorname{dom} f} \{\langle y, x \rangle - f(x)\},$$

which is always a closed and convex function.

For weakly convex functions $f \in \mathcal{F}_{\mu,\infty}$, with $\mu < 0$, one may define the subdifferential following [14, Proposition 6.3]. Let $\tilde{f}(x) \coloneqq f(x) - \mu \frac{\|x\|^2}{2}$, which is convex and has a well-defined subdifferential $\partial \tilde{f}(x)$, and let

$$\partial f(x) \coloneqq \{\tilde{g} + \mu x \,|\, \tilde{g} \in \partial \tilde{f}(x)\}.$$

If f is (weakly) convex, then the subdifferential notions coincide, i.e., $\partial f = \hat{\partial} f = \partial_L f$ [14, Proposition 6.2]. Moreover, if f is differentiable at x, then $\partial f(x) = \{\nabla f(x)\}$.

Definition 2.1.2 (Proximal mapping)**.** *The proximal mapping of a proper l.s.c. function $f\colon \mathbb{R}^d \to \bar{\mathbb{R}}$ with parameter $\gamma > 0$ is the set-valued operator*

$$\operatorname{prox}_{\gamma f}(x) \coloneqq \operatorname*{arg\,min}_{w \in \mathbb{R}^d} \{f(w) + \frac{1}{2\gamma}\|w - x\|^2\}. \tag{2.1}$$

The proximal mapping can be equivalently expressed in terms of the resolvent of ∂f as $\operatorname{prox}_{\gamma f}(x) = (\mathrm{id} + \gamma^{-1}\partial f)^{-1} x$. The first-order optimality conditions defining the proximal mapping, together with [98, Theorem 10.1, Exercise 8.8], imply

$$\gamma^{-1}(x - \bar{x}) \in \hat{\partial} f(\bar{x}), \quad \forall \bar{x} \in \operatorname{prox}_{\gamma f}(x).$$

2.2 Performance Estimation Methodology

The worst-case analysis framework, formalized by Nemirovski and Yudin [85, 83, 84], defines the performance of an algorithm over a function class by its maximum possible error within that class, based on black-box oracle information such as function values, gradient or proximal mapping values, Hessian, or higher-order derivatives. This perspective has been fundamental in understanding the limits of optimization methods and inspired the development of the celebrated accelerated gradient descent method of Nesterov [86]. For completeness, we note that while worst-case analysis is the standard benchmark and the framework adopted in this work, it can sometimes provide an overly pessimistic view, since the extreme instances defining the worst case may be rarely observed in practice.

Drori and Teboulle tackled the problem of finding the worst-case convergence rate of first-order optimization methods with access to a black-box oracle in a novel and systematic way in [44], introducing the performance estimation problem (PEP), a general tool for analyzing first-order methods. This framework was further refined and formalized by Taylor, Hendrickx and Glineur in [116, 113]. The key idea is to formulate the search for the most unfavorable behaviour of an algorithm within a given problem class as an optimization problem itself. By expressing this problem using inequalities that characterize the function class, PEP can often be reformulated as a convex semidefinite program (SDP), which can be solved efficiently.

Taylor et al. [116, 113] identified interpolation inequalities as necessary and sufficient conditions ensuring that the SDP formulation is tight, thereby providing a principled way to derive exact performance bounds. The PEP framework not only enables exact or tight convergence guarantees, but also provides a practical tool to compare algorithms, test the convergence of newly developed methods, and optimize their parameter settings based on worst-case behaviour. An important observation from this methodology is that any tight proof can be expressed as a linear combination of the interpolation inequalities describing the problem.

The PEP framework has since evolved and been applied to study many algorithms and to design new ones. For comprehensive views on PEP, one can check the following list of references [110, 111, 32, 1, 54, 19]. Notably, PEP has led to the development of the fastest known first-order methods in the worst-case sense for smooth functions: optimized gradient methods OGM [66] and OGM-G [65] for convex problems, and the triple momentum method [125] and the information-theoretic exact method [112] for strongly convex ones. Inexact oracle models can also be incorporated into the PEP framework, as shown in [11, 126]. An alternative line of work uses integral quadratic constraints (IQC) [75], while tight Lyapunov analysis for first-order methods can be automatically generated using some ideas in common with PEP [117, 123, 124].

We now describe the general PEP procedure, tailored to the gradient descent algorithm in Section 2.2.1. Consider an optimization problem over a function class $\mathcal{F}$ and a method $\mathcal{M}$. We seek to compute the worst-case behaviour of $\mathcal{M}$ when applied to a function in $\mathcal{F}$ (e.g., smooth/nonsmooth, convex/nonconvex). The worst-case behaviour is defined relative to an arbitrary starting point (possibly under constraints) and a specific performance criterion. Specifically, we are interested in computing the largest possible error in the result produced by the method.

Intuitively, the distance to the solution after performing N iterations of $\mathcal{M}$ is maximized over the class of functions $\mathcal{F}$. To ensure boundedness, we must specify the initial conditions and usually a characterization of the solution. A general form of the PEP problem is the following:

$$\begin{aligned} \underset{f\in\mathcal{F},(x_i)_{i=0,\dots,N},x_*}{\text{maximize}} \quad & \text{Performance Measure}\left(f,(x_i)_{i=0,\dots,N},x_*\right) \\ \text{subject to} \quad & \text{Iterates } x_i \text{ generated by method } \mathcal{M} \text{ from } x_0,\ i=1,\dots,N \\ & \text{Initial conditions on } x_0 \\ & \text{Optimality conditions on } x_*, \end{aligned} \tag{2.2}$$

where the decision variables are the function $f \in \mathcal{F}$ and the initial point x_0, while x_* denotes a minimizer.

Examples of performance measures include the optimality gap $f(x_N) - f_*$, distance to the optimum $\|x_N - x_*\|$, residual gradient $\|\nabla f(x_N)\|$ or best residual gradient (best stationarity measure) $\min_{0\leq i\leq N}\{\|\nabla f(x_i)\|\}$, where x_N is the final iterate generated by $\mathcal{M}$. Common initial conditions include bounds on the initial distance to optimum $\|x_0 - x_*\| \leq R$, initial optimality/function gap $f(x_0) - f_* \leq \Delta$ or initial stationarity residual $\|\nabla f(x_0)\| \leq R$, with constants $R, \Delta > 0$. If present, the optimality conditions provide a characterization of the minimizer, ranging from the trivial inequality $f(x_i) \geq f_*$ to tighter forms (see Lemma 3.6.1 for such a condition in the context of gradient descent).

Problem (2.2) is infinite dimensional because the maximization is taken over a class of functions. To obtain a finite formulation, one uses necessary and sufficient interpolating conditions for functions in $\mathcal{F}$. These conditions were first established in the context of PEP by Taylor et al. [116, 113]. In Section 2.2.2, we formally introduce these conditions, together with a short extension that covers weakly convex functions, and we also provide interpolation conditions characterizing prox-bounded functions.

Black-box model. We generally assume access to a first-order oracle represented by the set of triplets $\mathcal{T} \coloneqq \{(x_i, g_i, f_i) \cup (x_i, v_i, f_i)\}_{i\in\mathcal{I}} \subseteq \mathbb{R}^d \times \mathbb{R}^d \times \mathbb{R}$, where $\mathcal{I}$ is an index set. This oracle provides information about the iterates x_i, their gradients g_i and/or their proximal mapping values v_i, and the corresponding function values f_i.

2.2.1 Example: PEP setup for gradient descent

We now illustrate the PEP methodology for the case $\mathcal{M}$ = gradient descent, applied to functions $f \in \mathcal{F}_{\mu,L}$ (see Definition 2.1.1).

The problem in (2.3) below instantiates the PEP formulation (2.2), using the minimum gradient norm over the iterations as the *performance measure* and a bound on the initial optimality gap as the *initial condition*. Given curvature parameters μ and L, a bound $\Delta > 0$ characterizing the initial objective distance from the optimum, a number of iterations N and a fixed sequence of stepsizes γ_i, we consider:

$$\begin{aligned} \underset{f,(x_i)_{i=0,\dots,N},x_*}{\text{maximize}} \quad & \min_{0\le i\le N} \left\{ \|\nabla f(x_i)\|^2 \right\} \\ \text{subject to} \quad & x_{i+1} = x_i - \gamma_i \nabla f(x_i), \quad i \in \{0,\dots,N-1\} \\ & f \in \mathcal{F}_{\mu,L} \\ & f(x_0) - f_* \le \Delta. \end{aligned} \tag{2.3}$$

Because the maximization is taken over the entire class $\mathcal{F}_{\mu,L}$, the problem is infinite dimensional. To make it tractable, we relax it and show that the relaxed version yields the same optimal value (see Proposition 2.2.1 below). This relaxation step is key in the PEP methodology and relies on restricting the function to its values on the iterates by enforcing interpolation conditions.

For functions in $\mathcal{F}_{\mu,L}$, the corresponding interpolation conditions are stated in Theorem 2.2.1, formulated in the next section. Specifically, we have a set of inequalities $Q_{i,j} \ge 0$ holding for all $i, j \in \mathcal{I}$, with $i \neq j$, where $\mathcal{I} = \{0, \dots, N, *\}$ and:

$$Q_{i,j} \coloneqq f_i - f_j - \langle g_j, x_i - x_j \rangle - \frac{1}{2L}\|g_i - g_j\|^2 - \frac{\mu}{2L(L-\mu)}\|g_i - g_j - L(x_i - x_j)\|^2. \tag{2.4}$$

These inequalities encode a linear combination of function values and quadratic terms in the iterates x_i, x_j and gradients g_i, g_j.

Based on these interpolating conditions, the infinite dimensional PEP (2.3) is transformed into the finite dimensional PEP (2.5):

$$\begin{aligned}
\underset{\{(x_i,g_i,f_i)\}_{i\in\mathcal{I}}}{\text{maximize}} \quad & \min_{0\le i\le N} \left\{ \|g_i\|^2 \right\} \\
\text{subject to} \quad & \left\{(x_i, g_i, f_i)\right\}_{i\in\mathcal{I}} \text{ satisfy (2.4)} \\
& x_{i+1} = x_i - \gamma_i g_i, \quad i \in 0,\dots,N-1 \\
& f_0 - f_* \le \Delta.
\end{aligned} \tag{2.5}$$

Proposition 2.2.1 (Equivalence of infinite- and finite-dimensional PEPs)**.** *The optimization problems* (2.3) *and* (2.5) *have the same optimal value.*

Proof. Problem (2.5) is a relaxation of (2.3). It suffices to show that for any feasible solution $\bar{\mathcal{T}} = \left\{(\bar{x}_i, \bar{g}_i, \bar{f}_i)\right\}_{i\in\mathcal{I}}$ of (2.5), there exists a function $f \in \mathcal{F}_{\mu,L}$ such that $f(\bar{x}_i) = \bar{f}_i$, $\nabla f(\bar{x}_i) = \bar{g}_i$, $\forall i \in \mathcal{I}$, and $\min_{x\in\mathbb{R}^d} f(x) = f_*$. Since $\bar{\mathcal{T}}$ is a feasible point of problem (2.5), the assumptions of Theorem 2.2.1 are valid, hence such a function exists. □

The finite dimensional formulation (2.5) is still nonconvex and thus intractable. A tractable convex relaxation, known as the *SDP lifting*, is obtained by introducing a Gram matrix G to represent all inner products between pairs from the set containing both the iterates x_i and the gradients g_i. This allows rewriting the quadratic expressions in terms of linear combinations of entries of G. The gradient descent iterates are expressed as $x_i = x_0 + \sum_{k=0}^{i-1} \gamma_k g_k$, so that all x_i except x_0 vanish from the formulation. Then, letting $G \coloneqq \begin{bmatrix} g_0 & g_1 & \dots & g_N & x_0 - x_* \end{bmatrix}^\top \begin{bmatrix} g_0 & g_1 & \dots & g_N & x_0 - x_* \end{bmatrix}$, we obtain the following tractable semidefinite program:

$$\begin{aligned}
\underset{G,\,\ell,\,x_0,\,\{(g_i,f_i)\}_{i\in\mathcal{I}}}{\text{maximize}} \quad & \ell \\
\text{subject to} \quad & Q_{i,j} = f_i - f_j + \operatorname{tr}(A_{ij}G) \ge 0,\ i \ne j \\
& f_* - f_0 + \Delta \ge 0 \\
& G_{(i,i)} - \ell \ge 0, \quad i = 0,\dots,N \\
& G \succeq 0,
\end{aligned} \tag{primal-PEP}$$

where the matrices A_{ij} encode the interpolation conditions for the class $\mathcal{F}_{\mu,L}$. A detailed derivation of these matrices and the SDP reformulation can be found in [116, Sections 3.2-3.3]. This relaxation is exact when the dimension d of the function is large enough, namely when $d \geq N+2$ (the so-called "large-scale setting"). For lower dimensions, an additional nonconvex rank constraint may be introduced [116, Theorem 5].

The proofs of performance bounds are obtained by solving the dual of (primal-PEP). Let σ_i denote the Lagrange multipliers for the gradient norm constraints $G_{(i,i)} - \ell \geq 0$, $\alpha_{i,j} \geq 0$ the multipliers for the interpolation inequalities, and $\tau \geq 0$ the multiplier for the initial condition $f_* - f_0 + \Delta \geq 0$. The Lagrangian is given by

$$\mathcal{L} := \ell(1 - \sum_{i=0}^{N} \sigma_i) + \sum_{i=0}^{N} \sigma_i G_{(i,i)} - \tau(f_0 - f_* - \Delta) + \sum_{i,j\in\mathcal{I}} \alpha_{i,j} Q_{i,j}.$$

The dual formulation is obtained by maximizing over the primal variables and reads as:

$$\begin{aligned} \underset{\tau,\lambda_{i,j}\geq 0}{\text{minimize}} \quad & \tau\Delta \\ \text{subject to} \quad & \tau v_I - \sum_{i,j\in\mathcal{I}} \alpha_{i,j} v_F^{i,j} = 0 \\ & \operatorname{diag}(\sigma_i) + \sum_{i,j\in\mathcal{I}} \alpha_{i,j} A_{i,j} \preccurlyeq 0, \end{aligned} \tag{dual-PEP}$$

where v_I and $v_F^{i,j}$ are coefficient vectors defined by $\langle F, v_I\rangle = (f_0 - f_*)$ and $\langle F, v_F^{i,j}\rangle = (f_i - f_j)$, with $F = \begin{bmatrix} f_0 - f_* & \dots & f_N - f_* \end{bmatrix}$.

The prototype proof results from using weak duality, namely that the dual is larger than the Lagrangian value for any feasible primal variables F and G and feasible dual variables τ and $\alpha_{i,j}$:

$$\tau(f_0 - f_*) - \ell(1 - \sum_{i=0}^{N} \sigma_i) - \sum_{i=0}^{N} \sigma_i G_{(i,i)} - \sum_{i,j\in\mathcal{I}} \alpha_{i,j} Q_{i,j} \geq 0.$$

Using that $\ell = \min_i\{G_{(i,i)}\}$, the factor of ℓ must cancel. Then dividing the inequality by τ gives the standard form used to derive descent lemmas: (with an abuse of notation, we preserve σ_i and $\alpha_{i,j}$)

$$\underbrace{(f_0 - f_*)}_{\text{initial condition}} - \underbrace{\sum_{i=0}^{N} \sigma_i \|g_i\|^2}_{\text{performance metric}} - \underbrace{\sum_{i,j\in\mathcal{I}} \alpha_{i,j} Q_{i,j}}_{\text{interpolation conditions}} \geq 0. \tag{target inequality}$$

A worst-case guarantee is obtained by taking appropriate linear combinations of all constraints (interpolation inequalities, initial conditions, possible optimality conditions). These multipliers, $\alpha_{i,j}$ and σ_i, are a solution of the dual problem, and hence define the analytical proof of the bound. In this work, we sought principled ways of obtaining the multipliers $\alpha_{i,j}$, for instance by forming sums of squares from the interpolation conditions and enforcing certain cancellations consistent with the primal solution.

A more general description of the general proof methodology can be found in the tutorial by Goujaud et al. [55] or in the PhD thesis [116]. It is also noted there that strong duality holds when the iterates x_i are appropriately removed from the basis of G (as we have shown above), ensuring the existence of strictly feasible points F and $G \succ 0$.

Generally, PEP setup is limited to the cases where the performance measure, the initial conditions and the optimality conditions are expressed as linear combinations of the function values and the entries of the Gram matrix G (hence quadratic in x_i and g_i). Moreover, the number of iterations N is fixed a priori, and the considered methods are non-adaptive (i.e., with predetermined stepsizes).

Software tools. The primal and dual problems (primal-PEP)-(dual-PEP) can be solved with any standard SDP solver. Alternatively, the finite dimensional formulation in (2.5) can be handled directly using the MATLAB toolbox PESTO [115] or the Python package PEPit [56], both of which automatically convert the problem into an equivalent SDP. Moreover, [57] provides comprehensive and up-to-date documentation of PEP-based results.

2.2.2 Interpolation conditions

Interpolation conditions are a set of inequalities that tightly characterize whether a function belongs to a specific function class $\mathcal{F}$, based on a set of triplets $\mathcal{T} = \{(x_i, g_i/v_i, f_i)\}_{i\in\mathcal{I}}$, where x_i are the iterates, g_i the gradients, v_i the proximal mappings and f_i the function values. The interpolation problem itself, described in [96, 70], has its historical roots in Whitney's work on the extension of continuous functions [130, 131]. As demonstrated by Taylor et al. [116], interpolation conditions are a key tool for the precise convergence analysis required to *exactly* solve the infinite dimensional performance estimation problems (PEPs).

Interpolation conditions for $\mathcal{F}_{\mu,L}$-functions

In this section we present the interpolation inequalities for $\mathcal{F}_{\mu,L}$-functions, which are central to our derivations. For smooth (strongly) convex functions they were introduced by Taylor et al. [116, Theorem 4], and for smooth nonconvex ($\mu = -L$) in [113, Theorem 3.10]. We show in [100, Theorem 3.2] that these conditions also hold for weakly convex functions, i.e., for $\mu \in (-\infty, 0)$ with $\mu \neq L$. The proof follows similar arguments to the aforementioned works and is provided below.

Definition 2.2.1 ($\mathcal{F}_{\mu,L}$-interpolable set [116, Definition 2])**.** *The set* $\mathcal{T} \coloneqq \{(x_i, g_i, f_i)\}_{i\in\mathcal{I}}$ *is called* $\mathcal{F}_{\mu,L}$*-interpolable if and only if there exists a function* $f \in \mathcal{F}_{\mu,L}$ *such that* $\nabla f(x_i) = g_i$ *and* $f(x_i) = f_i$ *for all* $i \in \mathcal{I}$*.*

Lemma 2.2.1 ($\mathcal{F}_{\mu,L}$ upper & lower bounds)**.** *Let* $L \in \bar{\mathbb{R}}$ *and* $\mu \leq L$*, and let* $f \in \mathcal{F}_{\mu,L}$*. Then, for all* $x, y \in \mathbb{R}^d$ *and* $g \in \partial f(y)$*, the following holds:*

$$\tfrac{\mu}{2}\|x-y\|^2 \leq f(x) - f(y) - \langle g, x-y\rangle \leq \tfrac{L}{2}\|x-y\|^2. \tag{2.6}$$

Proof. Consider the *subgradient inequality* for convex functions,

$$h(x) \geq h(y) + \langle g, x-y\rangle\,,\ \forall x, y \in \mathbb{R}^d,$$

where $g \in \partial f(y)$. The two inequalities follow by setting $h = f - \frac{\mu}{2}\|\cdot\|^2$ and $h = \frac{L}{2}\|\cdot\|^2 - f$. □

Theorem 2.2.1 ($\mathcal{F}_{\mu,L}$-interpolation)**.** *Given an index set* $\mathcal{I}$*, consider the set of triplets* $\mathcal{T} \coloneqq \{(x_i, g_i, f_i)\}_{i\in\mathcal{I}} \subseteq \mathbb{R}^d \times \mathbb{R}^d \times \mathbb{R}$*. There exists* $f \in \mathcal{F}_{\mu,L}$*, with* $L \in \bar{\mathbb{R}}$ *and* $\mu \in (-\infty, L]$*, such that* $f(x_i) = f_i$ *and* $g_i \in \partial f(x_i)$ *for all* $i \in \mathcal{I}$*, if and only if the following the interpolation inequality* $Q_{i,j} \geq 0$ *holds for every pair of indices* (i,j)*, with* $i, j \in \mathcal{I}$*, where*

$$\begin{aligned} Q_{i,j} \coloneqq f_i - f_j - \langle g_j, x_i - x_j\rangle - \frac{1}{2L}\|g_i - g_j\|^2 \\ - \frac{\mu}{2L(L-\mu)}\|g_i - g_j - L(x_i - x_j)\|^2, \end{aligned} \tag{2.7}$$

when $L \neq \mu$*, while for* $L = \mu$ *it is*

$$Q_{i,j} \coloneqq f_i - f_j - \langle g_j, x_i - x_j\rangle - \frac{L}{2}\|g_i - g_j\|^2.$$

Proof. The proof follows the same steps as in [116, Theorem 4], with a minimal extension to accommodate $\mu < 0$. This extension impacts only the first step of the demonstration, known as minimal curvature subtraction, which involves shifting the lower curvature of the function to 0, effectively converting it into a convex function. This adjustment remains valid for $\mu < 0$, as shown below.

Let $\bar{f} \in \mathcal{F}_{0,L-\mu}$ be defined by $\bar{f}(x) \coloneqq f(x) - \frac{\mu}{2}\|x\|^2$, with $\partial\bar{f}(x) = \partial f(x) - \mu x$. By substituting f with $\bar{f}$ in the quadratic bounds inequality (2.6), we obtain:

$$0 \leq \bar{f}(x) - \bar{f}(y) - \langle \bar{g}_y, x-y\rangle \leq \frac{L-\mu}{2}\|x-y\|^2,$$

where $\bar{g}_y \in \partial\bar{f}(y)$. Hence, the set $\{(x_i, g_i, f_i)\}_{i\in\mathcal{I}}$ is $\mathcal{F}_{\mu,L}$-interpolable if and only if the set $\{(x_i, g_i - \mu x_i, f_i - \frac{\mu}{2}\|x_i\|^2)\}_{i\in\mathcal{I}}$ is $\mathcal{F}_{0,L-\mu}$-interpolable. The remainder of the proof follows identically to [116, Theorem 4]. □

Theorem 2.2.1 provides tighter conditions than [120, Theorem 2.2], which, to the best of our knowledge, is the first to explicitly include lower curvature information in weakly convex settings. The authors use it to derive the previous state-of-the-art convergence rates for the Douglas–Rachford splitting method; using Theorem 2.2.1 enables a tighter analysis, as we show in Chapter 5.

The terminology introduced in Definition 2.2.2 is used when describing the necessary and sufficient inequalities in the proofs.

Definition 2.2.2 (Distance-N interpolation conditions)**.** *Inequalities $Q_{i,j} \geq 0$ with $|i-j| = N$ are referred to as distance-N interpolation conditions.*

Corollary 2.2.1 ($\mathcal{F}_{\mu,\infty}$-interpolation)**.** *The set $\{(x_i, g_i, f_i)\}$ is $\mathcal{F}_{\mu,\infty}$-interpolable, with $\mu \in \mathbb{R}$, if and only if the conditions (2.7) hold for every pair of indices (i,j), with $i,j \in I$:*

$$f_i - f_j - \langle g_j, x_i - x_j\rangle \geq \frac{\mu}{2}\|x_i - x_j\|^2.$$

Proof. The result follows by taking the limit $L \to \infty$ in (2.7). □

Corollary 2.2.2 (Co-coercivity for weakly convex functions)**.** *Let $f \in \mathcal{F}_{\mu,L}$, with $\mu \in (-\infty, L)$. For all $x_i, x_j \in \mathbb{R}^d$, the following holds:*

$$0 \geq \langle \nabla f(x_i) - \nabla f(x_j) - L(x_i - x_j)\,,\, \nabla f(x_i) - \nabla f(x_j) - \mu(x_i - x_j)\rangle\,. \quad (2.8)$$

Proof. Inequality (2.8) is obtained by summing the inequalities (2.7) written for the pairs (x_i, x_j) and (x_j, x_i) and simplifying the result. □

Corollary 2.2.2 extends the co-coercivity property from [89, Theorem 2.1.12] to negative lower curvature values ($\mu < 0$). When $L + \mu > 0$, the class $\mathcal{F}_{\mu,L}$ fits within the framework of $(\frac{L\mu}{L+\mu}, \frac{1}{L+\mu})$-semimonotone operators (see, e.g., [47, Proposition 4.13]), which satisfy (2.8) by definition.

Interpolation conditions for prox-bounded functions

We now consider a more general class of functions. Specifically, assume that f is proper, l.s.c. and its proximal subproblems are well defined for the stepsizes of interest, i.e., prox-bounded (see Definition 2.2.3 below). Together with the interpolation conditions for $\mathcal{F}_{\mu,L}$-functions from Theorem 2.2.1, the conditions in this subsection form the foundation for deriving tight performance bounds for the Douglas–Rachford splitting method in the nonconvex setting analyzed in Chapter 5, which considers a smooth objective combined with a prox-bounded function.

Definition 2.2.3 (Prox-boundedness). *An l.s.c. function f is called γ-prox-bounded (or simply prox-bounded) if $f + \frac{1}{2\gamma}\|\cdot\|^2$ is bounded from below for some $\gamma > 0$.*

Definition 2.2.4 (Prox-boundedness threshold). *The prox-boundedness threshold of a function f is*

$$\gamma_f \coloneqq \sup\Big\{\gamma > 0\colon \inf_{x\in\mathbb{R}^d} \{f(x) + \frac{1}{2\gamma}\|x\|^2\} > -\infty\Big\}.$$

If f is bounded from below, then $\gamma_f = +\infty$.

With f being l.s.c., $\operatorname{prox}_{\gamma f}$ is nonempty and compact-valued on $\mathbb{R}^d$ for all $\gamma \in (0, \gamma_f)$ [98, Theorem 1.25].

Let $\{(x_i, w_i, f_i)\}_{i\in\mathcal{I}} \subset \mathbb{R}^d \times \mathbb{R}^d \times \mathbb{R}$ be a finite collection, where x_i are candidate proximal points, w_i are proximal centers, and f_i are prescribed function values. For prox-bounded functions, the following theorem characterizes when such finite data can be interpolated.

Theorem 2.2.2 (Prox-bounded interpolation). *Let $\gamma > 0$ and let $\{(x_i, w_i, f_i)\}_{i\in\mathcal{I}} \subset \mathbb{R}^d \times \mathbb{R}^d \times \mathbb{R}$, with $d \in \mathbb{N}$ and $\mathcal{I}$ being a finite index set. Define*

$$PB^w_{x,y} \coloneqq f(y) - f(x) + \frac{1}{2\gamma}\langle y - x, y + x - 2w\rangle. \tag{2.9}$$

Then there exists a proper l.s.c. γ-prox-bounded function $f\colon \mathbb{R}^d \to \bar{\mathbb{R}}$ such that $f_i = f(x_i)$ and $x_i \in \operatorname{prox}_{\gamma f}(w_i)$ for all $i \in \mathcal{I}$ if and only if $PB^{w_i}_{x_i,x_j} \geq 0$ for all $i, j \in \mathcal{I}$.

Proof. **Necessity.** If $x_i \in \operatorname{prox}_{\gamma f}(w_i)$, then by definition of the proximal mapping (see Definition 2.1.2), for all $x \in \mathbb{R}^d$,

$$f(x) + \frac{1}{2\gamma}\left\|x - w_i\right\|^2 \geq f(x_i) + \frac{1}{2\gamma}\left\|x_i - w_i\right\|^2.$$

Plugging in $x = x_j$ yields:

$$f(x_j) - f(x_i) + \frac{1}{2\gamma} \langle x_j - x_i \, , \, x_j + x_i - 2w_i \rangle \geq 0,$$

which is exactly $PB^{w_i}_{x_i,x_j} \geq 0$.

Sufficiency. Define

$$f(x) := \begin{cases} f_i, & x = x_i \text{ for some } i \in \mathcal{I}, \\ +\infty, & \text{otherwise.} \end{cases}$$

This function is proper and l.s.c., and it is γ-prox-bounded because

$$\inf_{x \in \mathbb{R}^d} \left\{ f(x) + \frac{1}{2\gamma} \|x\|^2 \right\} = \min_{i \in \mathcal{I}} \left\{ f_i + \frac{1}{2\gamma} \|x_i\|^2 \right\} > -\infty.$$

For a fixed $i \in \mathcal{I}$, since $f(x) = +\infty$ for all $x \notin \{x_j : j \in \mathcal{I}\}$, minimizing the proximal objective reduces to a finite comparison over $\{x_j\}_{j \in \mathcal{I}}$, i.e.,

$$\operatorname*{arg\,min}_{x \in \mathbb{R}^d} \left\{ f(x) + \frac{1}{2\gamma} \|x - w_i\|^2 \right\} = \operatorname*{arg\,min}_{j \in \mathcal{I}} \left\{ f_j + \frac{1}{2\gamma} \|x_j - w_i\|^2 \right\}.$$

Since $PB^{w_i}_{x_i,x_j} \geq 0$, x_i attains the minimum, i.e., $x_i \in \operatorname{prox}_{\gamma f}(w_i)$. □

2.2.3 PEP-type results

There exist several levels of PEP-related results, which can be organized as follows:

i) **Numerical worst-case values (primal-PEP solutions).** These correspond to purely numerical evaluation of the worst-case performance. Such results can be compared with analytical formulas to assess whether there is potential for improvement of known bounds.

ii) **Analytical worst-case performance bounds.** From the numerical results, one identifies the analytical expression of the worst-case value, thereby providing an interpretable and verifiable formula.

iii) **Formally proved bounds (dual-PEP solutions).** From the dual PEP solutions, one can identify the structure of the proofs and the corresponding multipliers of the interpolation inequalities involved. In some cases, this approach can be more straightforward than directly inferring the worst-case value and may provide a path to obtaining it analytically. Then, typically one must either construct a proof based on a sum of squares representation or, alternatively, verify that a corresponding slack matrix is positive semidefinite. This verification step often represents the most technically challenging part of the analysis.

iv) **Tightness and worst-case functions (primal-PEP solutions).** Constructing explicit functions that achieve the worst-case behaviour enables a formal verification of the tightness of the obtained bounds. The primal-PEP solution provides triplets that can often be interpolated by functions within the considered class. These triplets are typically easier to identify than dual multipliers and are often one-dimensional, allowing direct construction of a worst-case function. However, multidimensional worst-case examples also exist in the PEP literature, including some presented in this thesis for gradient descent. In such cases, one may instead verify that the PEP triplets are interpolable, a task that can be particularly challenging, as discussed in Chapter 3.

All our analyses provide analytical worst-case values on the entire feasible domain of the considered methods and function classes. We provide complete proofs of upper bounds for gradient descent (Chapter 3) and the Douglas–Rachford splitting (Chapter 5), and almost complete (up to a certain subdomain) for the difference-of-convex algorithm (Chapter 4). Moreover, for gradient descent we formally prove that these bounds are exact.

Chapter 3

Gradient descent: a complete worst-case analysis

Based on

[100] Teodor Rotaru, François Glineur, and Panagiotis Patrinos. Exact worst-case convergence rates of gradient descent: a complete analysis for all constant stepsizes over nonconvex and convex functions. *arXiv preprint arXiv:2406.17506*, 2024
(accepted for publication in the journal Mathematical Programming [99], doi:10.1007/s10107-025-02313-1)

3.1 Introduction

Tight convergence rates for the canonical first-order optimization algorithm, i.e., the gradient descent method, applied to smooth functions are demonstrated. We analyse functions with arbitrary upper and lower curvatures, covering weakly convex, convex and strongly convex cases. For all constant stepsizes ensuring descent, we derive tight bounds and establish convergence rates parameterized by curvature bounds, stepsize, and iteration count.

The key tool enabling us to obtain a complete characterization is the analysis of interconnection of iterations not only for consecutive steps, but also for the ones lying at distance two, unlike classical convergence analysis, e.g., [89, 23, 17]. Standard PEP formulations typically rely on non-consecutive iterates to characterize worst-case performance. However, incorporating all resulting dependencies generally leads to substantially harder proofs; consequently, in practice, relatively few analyses leverage more than consecutive iterates.

Our proof strategy slightly differs from existing work on PEP (e.g., [65]), which connects consecutive iterates to the final one, as well as from the recent line of work on stepsize-based acceleration (e.g., via silver stepsizes), which exploits only a subset of inter-iterate connections. In contrast, our approach enables an analysis that is decoupled from the final iterate, relying instead on standard descent lemmas to establish performance bounds. To the best of our knowledge, results covering the full range of stepsizes had previously been, at most, conjectured, even in the apparently simpler scenario of convex objectives.

3.1.1 Notations and Definitions

We consider the unconstrained optimization problem

$$\underset{x\in\mathbb{R}^d}{\text{minimize}}\, f(x),$$

where f belongs to the class of functions $\mathcal{F}_{\mu,L}$ (recall Definition 2.1.1). Recall that, based on the sign of μ, the function f is (i) weakly convex (or hypoconvex) for $\mu < 0$, (ii) convex for $\mu = 0$ or (iii) strongly convex for $\mu > 0$. Moreover, its Lipschitz gradient constant is $\max\{-\mu, L\}$. Note that the class of smooth weakly convex functions includes the smooth not-necessarily convex ones under the condition $\mu = -L$. The trivial case $\mu = L$ is not explicitly addressed. It leads to a simple quadratic problem, whose Hessian with all eigenvalues equal to L and for which the answers are easy to derive (for a one-dimensional example, see Proposition 3.6.1).

Gradient descent. Starting from some initial point $x_0 \in \mathbb{R}^d$, we consider N iterations of gradient descent, as in Algorithm 1. We restrict the analysis to fixed stepsizes $\gamma_i \in (0, \frac{2}{L})$, ensuring a decrease in the function value after a single iteration. We assume the function f to be bounded from below and denote $f_* \coloneqq \inf_x f(x) > -\infty$. In the nonconvex case, where first-order methods cannot guarantee convergence to a global minimizer, analysis focuses on the stationarity measure $\min_{0\le i\le N}\{\|\nabla f(x_i)\|^2\}$, typically bounded in terms of the initial optimality global gap $f(x_0) - f_*$, or, when f_* is unknown, by the observable decrease $f(x_0) - f(x_N)$.

Algorithm 1: Gradient descent algorithm (GD)

Data: $f \in \mathcal{F}_{\mu,L}$, with $L > 0$ and $\mu \in (-\infty, L)$; $N \geq 1$ iterations starting from $x_0 \in \mathbb{R}^d$

for $i = 0, \ldots, N$ **do**

$$x_{i+1} = x_i - \gamma_i \nabla f(x_i) \tag{GD}$$

Result: Best iterate x_i with $i = \arg\min_{0 \leq i \leq N} \{\|\nabla f(x_i)\|^2\}$

3.1.2 Prior work

The conventional analysis of gradient descent relies on a constant stepsize schedule and provides non-tight rates (for example, see the textbooks [89, 17, 23, 25, 15]).

Convergence rates of first-order methods applied to smooth convex functions have been extensively examined using PEP. For gradient descent, in [44] tight rates are proved for smooth convex functions with stepsizes $\gamma L \in (0, 1]$ and conjectured for $\gamma L \in (1, 2)$; exact rates for smooth strongly convex functions are conjectured for $\gamma L \in (0, 2)$ [116] and proved when employing line-search [37]. Optimized first-order methods can be derived within the PEP framework, e.g., in [65] for an optimal algorithm to decrease the gradient norm of smooth convex functions. A recent breakthrough by Teboulle [118] establishes tight convergence rates for smooth convex functions using stepsizes $\gamma L \in (1, \frac{3}{2}]$.

For smooth nonconvex functions, sometimes encountered in fields like machine learning, there is a notable gap in comprehending convergence rates. Carmon et al. establish lower bounds for finding stationary points in [26, Theorem 2], demonstrating that gradient descent is worst-case optimal for functions with only a Lipschitz gradient. In a related study, Cartis et al. present tight lower bounds for steepest descent [28]. Notably, by additionally assuming Lipschitz continuity of the second derivatives, Carmon et al. [27] derive accelerated gradient descent for weakly-convex functions. Within the PEP framework, Taylor examines the convergence rates of gradient descent applied to smooth, not-necessarily convex functions, i.e., $\mu = -L$, with a constant stepsize $\frac{1}{L}$ [110, p. 190]. Drori and Shamir extend this result to the case of variable, yet predetermined, stepsizes below $\frac{1}{L}$ [43, Corollary 1] (using scaling arguments, this extension could also be derived from the result in [110, p. 190]). Abbaszadehpeivasti et al. further improve upon these findings [2], achieving state of the art for stepsizes up to $\frac{\sqrt{3}}{L}$. Typically, these aforementioned results solely rely on exploiting smoothness, neglecting the weak convexity parameter μ. Recent interest in weakly convex optimization has emerged, e.g., [36, 8].

Given our focus on constant stepsize schedule, our analysis is confined to the interval $\gamma L \in (0, 2)$. However, recent advancements address longer stepsizes and achieve improved asymptotic rates. Interested readers are directed to the works of Altschuler and Parrilo [5, 6], Das Gupta et al. [34] and Grimmer et al. [58, 59, 60]. Notably, similar to our approach, proving their rates requires connecting more than just consecutive iterations, deviating from classical analyses. Comparable demonstrations are provided in [65, 138].

Nesterov highlighted the significance of minimizing the gradient norm [87], asserting that $\|\nabla f(x)\|$ might serve better than other performance metrics for minimization purposes. Evaluating the gradient norm for initially bounded functions is a common approach in analyzing gradient-based methods for smooth nonconvex optimization, seen in works such as [43] for stochastic gradient descent and in [65] for the OGM-G algorithm optimizing the gradient norm for smooth convex functions. This convergence measure is suitable for applications where both the initial optimality gap and the gradient norm across iterations can be directly measured, unlike metrics involving the optimizer. We maintain consistency by utilizing this performance metric for both convex and strongly convex functions.

3.1.3 Chapter organization and main contributions

Our contributions can be summarized as follows, with a comprehensive comparison to the state of the art provided in Table 3.1:

1. We provide tight performance bounds for gradient descent with constant stepsizes $\gamma L \in (0, 2)$, proving their tightness through worst-case examples. Our analysis applies to weakly convex (Theorem 3.2.3), strongly convex (Theorem 3.2.2), and convex functions (Theorem 3.2.1), with a comprehensive summary in Table 3.2.

2. A performance bound for not-necessarily convex functions using variable stepsizes $\gamma_i L \in (0, \overline{\gamma L}_1(\kappa)]$[1] (Theorem 3.2.4), which we prove is tight for stepsizes below 1 and conjecture for the remaining range. This result generalizes the constant stepsize case for weakly convex functions (Theorem 3.2.3).

3. Leveraging our tight performance bounds, we propose optimized stepsize rules for gradient descent, namely: (i) best constant stepsize for convex, strongly

[1] The quantities $\overline{\gamma L}_k(\kappa)$ denote normalized stepsize thresholds that depend on the curvature ratio κ. They are monotonically increasing in the index $k \geq 0$ and are formally defined in Definition 3.2.2.

Table 3.1: State of the art vs contributions for gradient descent analysis. Legend: **Bold** = complete/proven results in this chapter. Abbreviations: w.c./c./s.c. = weakly convex / convex / strongly convex.

Criterion	**State of the art**	**This chapter**
Function class	Partial coverage: c./s.c. [44, 116]; w.c. mainly $\mu = -L$ slices [110, 43, 2]	**Complete**: w.c./c./s.c.; $f \in \mathcal{F}_{\mu,L}$
Constant stepsizes	c.: $\gamma L \in (0, 1]$ [44]; $\gamma L \in (1, \frac{3}{2}]$ [118]; s.c.: conjectured for $\gamma L \in (0, 2)$ [116]; w.c.: up to $\frac{\sqrt{3}}{L}$ [2]	**Full range**: $\gamma L \in (0, 2)$ for all curvature regimes
Performance metric	Mixed: function gap/distance for c./s.c. [44, 116]; gradient norm for w.c. [26]	**Unified**: $\min_i \|\nabla f(x_i)\|^2$ subject to $f(x_0) - f(x_N)$ or $f(x_0) - f_*$
Proved upper bounds	c.: $\gamma L \in (0, 1]$ [44], $\gamma L \in (1, \frac{3}{2}]$ [118]; s.c.: $\gamma L \in (0, 2)$ for some criteria [116, 113], line-search [37]; w.c. ($\mu = -L$): $\gamma_i L \leq 1$ [43], $\gamma L \leq \sqrt{3}$ [2]	**All**: $\gamma L \in (0, 2)$ for w.c./c./s.c.
Proved tightness (lower bounds)	c.: $\gamma L \in (0, \frac{3}{2}]$ [44, 118]; s.c.: $\gamma L \in (0, 2)$ for various criteria [116, 113], line-search [37]	**Complete**: worst-case examples for all bounds
Proof technique	Standard PEP with consecutive iterates or final-iterate coupling [65]	**Distance-2 coupling**: three-iterate interconnections
Stepsize design	Limited rules; predetermined sequences $\gamma_i \leq 1/L$ for w.c. ($\mu = -L$) [43]; optimized methods for c. [65]; stepsize-based acceleration [5, 6, 58, 59, 60, 34]	**Focus on constant stepsizes**: iteration-count optimal + horizon-free dynamic policy

convex and weakly convex functions (Propositions 3.3.1 to 3.3.3) for a given number of iterations; (ii) a dynamic sequence of stepsizes, independent of the iteration count N, achieving a superior guarantee compared to any constant stepsize policy (Theorem 3.3.1).

The aforementioned works typically address only the convex case ($\mu = 0$), the strongly convex case ($\mu \in (0, L)$), the Lipschitz smooth case ($\mu = -L$), or the setting without lower curvature information ($\mu = -\infty$) (see, e.g., [89, §1.2.3]). Our framework unifies these regimes and yields a continuous spectrum of performance guarantees interpolating between them. We recall that the

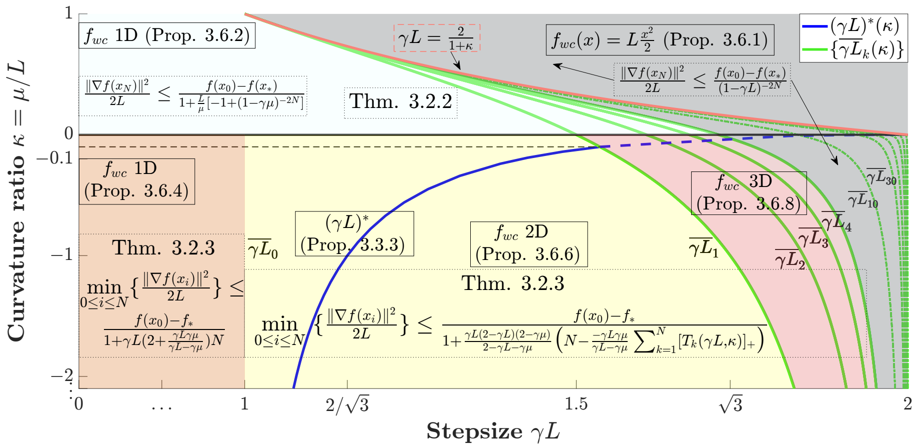

Figure 3.1: Regimes outlined in Theorem 3.2.1 (smooth convex), Theorem 3.2.2 (smooth strongly convex) and Theorem 3.2.3 (smooth weakly convex) with respect to the constant stepsize γL and curvature ratio κ, for $N = 4$ iterations. State-of-the-art *proven* tight convergence rates are limited to the cases $\kappa = -\infty$ [89], $\kappa = -1$ [2] and $\kappa = 0$ [44, 118], with stepsizes lower than $\overline{\gamma L}_1(\kappa)$. Above this threshold, a series of N regimes parameterized by $k = 1, \dots, N$ exist and they are sublinear for non-strongly convex functions. Stepsize thresholds converge to $\frac{2}{1+[\kappa]_+}$ as $N \to \infty$, as illustrated for $k = \{10, 30, 50, 100, 200, 1000\}$. The optimal constant stepsize $(\gamma L)^*(\kappa)$ for weakly convex functions is independent of N for $\kappa \lesssim -0.1$, otherwise it is only asymptotically valid (as marked with *blue* dashed line). We use inequalities connecting three consecutive iterations to prove the regimes corresponding to stepsizes $\gamma L \in (\overline{\gamma L}_1, \frac{2}{1+[\kappa]_+})$. The worst-case function (f_{wc}) dimensionality and the tightness' proof references are illustrated.

performance metric for (strongly) convex functions usually differs, e.g., in [44, Conjecture 1], [116, Conjecture 3] and [114, Theorem 2.1].

The gist in proving tight rates on the full interval of stepsizes is its partition into a collection of intervals delimited by stepsize thresholds dependent on the curvature ratio κ, each of them requiring distinct analysis. Figure 3.1 illustrates this partitioning for curvature ratio and (normalized) stepsizes, along with their corresponding rates for $N = 4$ iterations.

The remainder of this chapter is organized as follows. Section 3.2 contains the formal statements of our main theoretical findings. Section 3.4 provides preliminary results required for the proofs. The complete proofs for the upper convergence bounds are detailed in Section 3.5, while Section 3.6 presents the worst-case examples that establish their tightness. To facilitate reproducibility, we provide a supplementary GitHub repository[2] containing MATLAB scripts for numerical verification and symbolic validation of the primary algebraic manipulations.

3.2 Performance bounds

In Sections 3.2.1 to 3.2.3 we provide the exact performance bounds for constant stepsize schedules. Table 3.2 summarizes the different results involved in their proofs. In Section 3.2.4 we give the tight bounds for smooth not-necessarily convex functions when using variable stepsizes, upper bounded by some threshold. Section 3.2.5 describes the challenging stepsize thresholds defining the regimes for constant stepsizes. Section 3.3 provides optimal stepsize rules with respect to the theoretical worst-case scenarios: best constant stepsize policies in Section 3.3.1 and a dynamic stepsize schedule with better worst-case guarantees in Section 3.3.2.

The convergence rates are influenced by: (i) curvature ratio $\kappa = \frac{\mu}{L}$, (ii) the (normalized) fixed stepsizes $\gamma_i L \in (0, 2)$, and (iii) the number of iterations N. Given a *fixed budget* (number of iterations N), the rates show stepsize dependent accuracies in finding a stationary point, of the type

$$\frac{1}{2L} \min_{0 \le i \le N} \left\{ \|\nabla f(x_i)\|^2 \right\} \le \frac{f(x_0) - f_*}{1 + \tau_N(\kappa, (\gamma_i L)_{i=0}^{N-1})}, \qquad (\mathcal{P}_*)$$

where $\tau_N(\cdot)$ is some function we determine through our analysis.

Remark 3.2.1 (Removing f_* in analysis)**.** *To be consistent with standard performance bounds in the literature, our final theorems use the optimality gap $f(x_0) - f_*$, as in ($\mathcal{P}_*$). We derive these results by providing a bound on the gap $f(x_0) - f(x_N)$, which is a measure directly available from the first-order oracle:*

$$\frac{1}{2L} \min_{0 \le i \le N} \left\{ \|\nabla f(x_i)\|^2 \right\} \le \frac{f(x_0) - f(x_N)}{\tau_N(\kappa, (\gamma_i L)_{i=0}^{N-1})}. \qquad (\mathcal{P}_N)$$

The two bounds differ only by a "+1" term in the denominator. In nonconvex settings, e.g., multimodal functions, possessing multiple local optima, ($\mathcal{P}_N$) can be more effective for characterizing local convergence as it avoids the distortion of the numerator from a distant global optimum f_* with a much lower value.

[2] GitHub repository: https://github.com/teo2605/GD_tight_rates

Table 3.2: Overview of performance bounds proofs and their tightness for all regimes with constant stepsizes (the thresholds $\overline{\gamma L}_k$, formally defined in Section 3.2.5, increase with k). The demonstrations use interpolation inequalities connecting consecutive iterations (distance-1) or, additionally, the ones situated at distance-2. Tightness is shown by constructing a worst-case function or by proving its existence via interpolating triplets.

Class	Stepsizes	Key result	Inequalities	Tightness proof	Worst-case example
convex $\mu = 0$ Theorem 3.2.1	$(0, \frac{3}{2}]$	Lemma 3.5.1 & Lemma 3.5.2	distance-1	Proposition 3.6.3	1D function
	$(\frac{3}{2}, \overline{\gamma L}_N)$	Corollary 3.5.1	distance-1&2		
	$[\overline{\gamma L}_N, 2)$	Lemma 3.5.7	distance-1&2	Proposition 3.6.1	
strongly convex $\mu > 0$ Theorem 3.2.2	$(0, \overline{\gamma L}_1]$	Lemma 3.5.4	distance-1	Proposition 3.6.2	1D function
	$[\overline{\gamma L}_1, \overline{\gamma L}_{N-1})$	Lemma 3.5.4 & Lemma 3.5.8	distance-1&2		
	$[\overline{\gamma L}_{N-1}, \overline{\gamma L}_N)$	Lemma 3.5.8	distance-1&2		
	$[\overline{\gamma L}_N, \frac{2}{1+\kappa})$	Lemma 3.5.7	distance-1&2	Proposition 3.6.1	
	$[\frac{2}{1+\kappa}, 2)$	Lemma 3.5.3	distance-1		
weakly convex $\mu < 0$ Theorem 3.2.3	$(0, 1]$	Lemma 3.5.1	distance-1	Proposition 3.6.4	1D function
	$[1, \overline{\gamma L}_1]$	Lemma 3.5.2	distance-1	Proposition 3.6.6	2D triplets
	$[\overline{\gamma L}_1, \overline{\gamma L}_{N-1})$	Lemma 3.5.2 & Lemma 3.5.6	distance-1&2	Proposition 3.6.8	3D triplets
	$[\overline{\gamma L}_{N-1}, \overline{\gamma L}_N)$	Lemma 3.5.6	distance-1&2	Proposition 3.6.7	2D function
	$[\overline{\gamma L}_N, 2)$	Lemma 3.5.7	distance-1&2	Proposition 3.6.1	1D function

3.2.1 Performance bounds for convex functions

In Theorem 3.2.1 we establish the exact performance bound for smooth convex functions. This result arises as a limiting case of Theorem 3.2.2 and Theorem 3.2.3. Since the smooth convex setting is the most commonly encountered regime in analyses of gradient descent, we present it as a theorem.

Theorem 3.2.1 (Exact performance bound for convex functions)**.** *Let $f \in \mathcal{F}_{0,L}$ and consider N iterations of (GD) with constant stepsizes $\gamma L \in (0, 2)$ starting from x_0. Then the following bound holds:*

$$\frac{1}{2L}\|\nabla f(x_N)\|^2 \leq \frac{f(x_0) - f_*}{1 + \gamma L \min\left\{2N\,,\ \frac{-1+(1-\gamma L)^{-2N}}{\gamma L}\right\}}. \tag{3.1}$$

We prove Theorem 3.2.1 in Section 3.5.6. In Proposition 3.6.3, we provide a worst-case function attaining the performance bound (3.1).

To our knowledge, this result establishes the first *proved* tight rate in the convex case for the full range of constant stepsizes. The denominator in (3.1) is similar to the numerically conjectured one from [44, Conjecture 3.1], even though for a different performance metric, with $x_* \in \arg\min_x f(x)$:

$$f(x_N) - f_* \leq \frac{L}{2} \frac{\|x_0 - x_*\|^2}{1 + \gamma L \min\left\{2N\,,\, \frac{-1+(1-\gamma L)^{-2N}}{\gamma L}\right\}}. \tag{3.2}$$

Since the denominator is the same, both metrics are optimized by the same optimal constant stepsize (see Proposition 3.3.1).

3.2.2 Performance bounds for strongly convex functions

In Theorem 3.2.2 we derive the exact performance bound for smooth strongly convex functions.

Theorem 3.2.2 (Exact performance bound for strongly convex functions)**.** *Let $f \in \mathcal{F}_{\mu,L}$, with $\mu \in (0, L)$, and consider N iterations of (GD) with stepsizes $\gamma L \in (0, 2)$ starting from x_0. Then the following bound holds:*

$$\frac{1}{2L}\|\nabla f(x_N)\|^2 \leq \frac{f(x_0) - f_*}{1 + \gamma L \min\left\{\frac{-1+(1-\gamma\mu)^{-2N}}{\gamma\mu}\,,\, \frac{-1+(1-\gamma L)^{-2N}}{\gamma L}\right\}}. \tag{3.3}$$

We prove Theorem 3.2.2 in Section 3.5.5. In Proposition 3.6.2, we provide a worst-case function that reaches the performance bound (3.3).

Tight performance bounds for measures different from ours are given in [114, Table 2][3]. For the metric analysed in this chapter, however, only the rate corresponding to the stepsize $\gamma = \frac{1}{L}$ is conjectured there. In [113], the tight contraction factor $\max\{\gamma L - 1, 1 - \gamma\mu\}$ is derived and is optimized by choosing $\gamma = \frac{2}{L+\mu}$. Notably, this stepsize choice is suboptimal for the metric considered in our work (see Proposition 3.3.2) and is also outperformed by the dynamic stepsize procedure proposed in Section 3.3.2.

Subsequent to the publication of our work on *arXiv* [100], an alternative proof of Theorem 3.2.2 was provided in [68]. In the same paper, exploiting the *H-duality* phenomenon [67], the author provides the first proof for the numerically

[3]The work in [114] provides tight performance bounds for the proximal gradient method, which generalizes gradient descent.

conjectured rate from [116, Conjecture 2], which in our notation writes as:

$$f(x_N) - f_* \le \frac{L}{2} \frac{\|x_0 - x_*\|^2}{1 + \gamma L \min\left\{\frac{-1+(1-\gamma\mu)^{-2N}}{\gamma\mu}, \frac{-1+(1-\gamma L)^{-2N}}{\gamma L}\right\}}, \tag{3.4}$$

where $x_* = \arg\min_x f(x)$. Since the denominator is the same, both metrics are optimized by the same optimal constant stepsize (see Proposition 3.3.2).

Letting $\mu \searrow 0$, the rate for convex functions from Theorem 3.2.1 is recovered.

3.2.3 Performance bounds for weakly convex functions

In this section we derive and discuss performance bounds for smooth weakly convex functions, when using constant stepsizes. We begin by stating the bound in Theorem 3.2.3.

Theorem 3.2.3 (Exact performance bound for weakly convex functions)**.** *Let $f \in \mathcal{F}_{\mu,L}$, with $\mu < 0$, and consider N iterations of* (GD) *with constant stepsizes $\gamma L \in (0, 2)$ starting from x_0. Then the following bound holds:*

$$\frac{1}{2L} \min_{0\le i\le N} \left\{\|\nabla f(x_i)\|^2\right\} \le \frac{f(x_0) - f_*}{1 + \gamma L \min\left\{P_N(\gamma L, \gamma\mu), \frac{-1+(1-\gamma L)^{-2N}}{\gamma L}\right\}}, \tag{3.5}$$

with $P_N(\ell, u)$ defined as:

$$P_N(\ell, u) := p(\ell, u)\left[N - \frac{-\ell u}{\ell - u} \sum_{k=1}^{N} \left[\frac{-1 + (1-u)^{-2k}}{u} - \frac{-1 + (1-\ell)^{-2k}}{\ell}\right]_+\right] \tag{3.6}$$

where

$$p(\ell, u) := 2 - \frac{-\ell u}{1 - u - |1-\ell|} = \begin{cases} 2 - \frac{-\ell u}{\ell - u}, & \ell \in (0, 1]; \\ \frac{(2-\ell)(2-u)}{2-\ell-u}, & \ell \in [1, 2). \end{cases} \tag{3.7}$$

We prove Theorem 3.2.3 in Section 3.5.4. We then establish the tightness of these rates over the entire stepsize interval, which we partition into several subintervals. Explicit worst-case functions are provided for $\gamma L \in (0, 1]$ (one-dimensional, Proposition 3.6.4), $\gamma L \in [\overline{\gamma L}_{N-1}(\kappa), \overline{\gamma L}_N(\kappa))$ (two-dimensional, Proposition 3.6.7) and $\gamma L \in [\overline{\gamma L}_N(\kappa), 2)$ (one-dimensional, Proposition 3.6.1). For the remaining intervals, the worst-case behaviour is captured by interpolating triplets: $\gamma L \in (1, \overline{\gamma L}_1(\kappa)]$ (two-dimensional, Proposition 3.6.6) and $\gamma L \in [\overline{\gamma L}_1(\kappa), \overline{\gamma L}_N(\kappa))$ (three-dimensional, Proposition 3.6.8).

The rate's denominator combines a set of N sublinear regimes with a linear regime also observed for convex and strongly convex functions (Theorem 3.2.1 and Theorem 3.2.2), holding for stepsizes close to 2. For $\gamma L \in (0,1]$, each term in the sum appearing in (3.6) is equal to zero. With $\gamma L \in (1,2)$, only the terms up to some index k become positive, while the remaining $(N-k)$ terms are negative. The expression is continuous with respect to the stepsize. A new term in the sum becomes nonnegative when evaluated at specific stepsize thresholds $\overline{\gamma L}_k$ rigorously defined in Section 3.2.5 and satisfying the condition:

$$\frac{-1+(1-\kappa\overline{\gamma L}_k)^{-2k}}{\kappa\overline{\gamma L}_k} - \frac{-1+(1-\overline{\gamma L}_k)^{-2k}}{\overline{\gamma L}_k} = 0. \tag{3.8}$$

Using the identity

$$\sum_{j=1}^{k} \frac{-1+(1-u)^{-2j}}{u} - \frac{-1+(1-l)^{-2j}}{l} =$$

$$\frac{-1+(1-u)^{-2k}}{1-(1-u)^2} - \frac{-1+(1-l)^{-2k}}{1-(1-l)^2} + \Big(\frac{1}{l} - \frac{1}{u}\Big)k,$$

the expression of $P_N(\gamma L, \gamma\mu)$ from (3.6) can be equivalently written as

$$P_N(\gamma L, \gamma\mu) = p(\gamma L, \gamma\mu) \min_{0\le k\le N} \left\{ \frac{\frac{-1+(1-\gamma L)^{-2k}}{\gamma L[1-(1-\gamma L)^2]} - \frac{-1+(1-\gamma\mu)^{-2k}}{\gamma\mu[1-(1-\gamma\mu)^2]}}{\frac{1}{\gamma L} - \frac{1}{\gamma\mu}} + N - k \right\}.$$

For stepsize thresholds satisfying (3.8), it reduces to summing an exponential term with a linear one:

$$P_N(\overline{\gamma L}_k, \kappa\overline{\gamma L}_k) = \frac{(2-\overline{\gamma L}_k)(2-\kappa\overline{\gamma L}_k)}{2-\overline{\gamma L}_k - \kappa\overline{\gamma L}_k}(N-k) + \frac{-1+(1-\overline{\gamma L}_k)^{-2k}}{\overline{\gamma L}_k}.$$

When reaching the maximum threshold $\overline{\gamma L}_N$ corresponding to $k = N$, $P_N(\gamma L, \gamma\mu)$ becomes $\frac{-1+(1-\overline{\gamma L}_N)^{-2N}}{\overline{\gamma L}_N}$ and the rate shifts to the regime linear in $(1-\overline{\gamma L}_N)$. Summarizing, the expression of $P_N(\gamma L, \gamma\mu)$ can be expanded as follows:

$$\begin{cases} \big(2 + \frac{\gamma\mu\gamma L}{\gamma L - \gamma\mu}\big)N, & \gamma L \in (0,1]; \\ \frac{(2-\gamma L)(2-\gamma\mu)}{2-\gamma L-\gamma\mu}(N-k) + \frac{\frac{-1+(1-\gamma L)^{-2k}}{\gamma L[1-(1-\gamma L)^2]} - \frac{-1+(1-\gamma\mu)^{-2k}}{\gamma\mu[1-(1-\gamma\mu)^2]}}{\frac{1}{1-(1-\gamma L)^2} - \frac{1}{1-(1-\gamma\mu)^2}}, & \gamma L \in \big[\overline{\gamma L}_k, \overline{\gamma L}_{k+1}\big) \\ & k = 0, \dots, N-1; \\ \frac{-1+(1-\gamma L)^{-2N}}{\gamma L}, & \gamma L = \overline{\gamma L}_N. \end{cases}$$

By fixing stepsize γL and varying N, another perspective on convergence rates emerges: determining the maximum number of iterations to reach a desired accuracy with a *given stepsize*. Let $\bar{N}$ be the largest index k corresponding to a strictly positive term in expression of $P_N(\gamma L, \gamma\mu)$ from (3.6). Then for $N \leq \bar{N}$ the rate is linear in $(1-\gamma L)$, whereas for $N \geq \bar{N}+1$ it shifts to a sublinear regime.

We end this subsection with Corollary 3.2.1 addressing the standard smooth and not-necessarily convex case, namely $\mu = -L$, extending the results from [2] holding for constant stepsizes up to $\frac{\sqrt{3}}{L}$.

Corollary 3.2.1 (Exact performance bound under smoothness only)**.** *Let $f \in \mathcal{F}_{-L,L}$ and consider N iterations of* (GD) *with constant stepsizes $\gamma L \in (0,2)$ starting from x_0. Then the following bound holds:*

$$\frac{1}{2L} \min_{0\leq i\leq N} \left\{\|\nabla f(x_i)\|^2\right\} \leq \frac{f(x_0) - f_*}{1 + \gamma L \min\left\{P_N(\gamma L)\,,\ \frac{-1+(1-\gamma L)^{-2N}}{\gamma L}\right\}},$$

with $P_N(\ell)$ defined as:

$$P_N(\ell) := p(\ell)\Big(N - \frac{1}{2}\sum_{k=1}^{N}\Big[2 - (1+\ell)^{-2k} - (1-\ell)^{-2k}\Big]_+\Big)$$

where

$$p(\ell) := \frac{4 - \ell\max\{1,\ell\}}{2} = \begin{cases} \frac{4-\ell}{2}, & \ell \in (0,1]; \\ \frac{4-\ell^2}{2}, & \ell \in [1,2). \end{cases}$$

3.2.4 Performance bounds for non-strongly convex functions and variable stepsizes

In this section, we derive and discuss performance bounds for smooth weakly convex and convex functions under variable (but predefined) stepsizes, up to a well-defined threshold. In Theorem 3.2.4, we establish the corresponding guarantee for $\mu \leq 0$, and Corollary 3.2.2 specializes this guarantee to the convex setting (i.e., $\mu = 0$).

Theorem 3.2.4 (Exact performance bound for non-strongly convex functions and variables stepsizes up to first threshold)**.** *Let $f \in \mathcal{F}_{\mu,L}$, with $\mu \in (-\infty, 0]$, and consider N iterations of* (GD) *starting from x_0 with stepsizes $\gamma_i L \in (0, \overline{\gamma L}_1(\kappa)]$, $i = 0,\dots,N-1$, where $\overline{\gamma L}_1(\kappa) := \frac{3}{1+\kappa+\sqrt{1-\kappa+\kappa^2}} \in [\frac{3}{2}, 2)$. Then the following bound holds, where $p(\gamma L, \gamma\mu)$ is defined in* (3.7)*:*

$$\frac{1}{2L} \min_{0\leq i\leq N} \left\{\|\nabla f(x_i)\|^2\right\} \leq \frac{f(x_0) - f_*}{1 + \sum_{i=0}^{N-1} \gamma_i L\, p(\gamma_i L, \gamma_i \mu)}. \tag{3.9}$$

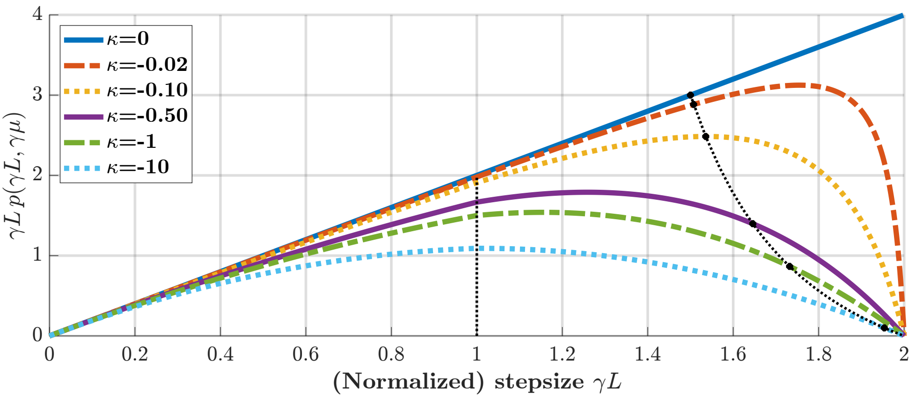

Figure 3.2: The dominant term $\gamma L\, p(\gamma L, \gamma\mu)$ (see (3.7)) in the bounds for weakly convex functions for several curvature ratios. $\gamma L = 1$ delimits the two main sublinear regimes, while the one-step analysis is tight up to the threshold $\overline{\gamma L}_1(\kappa)$, marked in black dots, above which intervene transient contributions (see $P_N(\gamma L, \gamma\mu)$ in (3.6)).

Theorem 3.2.4 is demonstrated in Section 3.5.4. The bound in (3.9) is proved to be tight for $\gamma_i L \in (0, 1]$ in Proposition 3.6.4. For the range $\gamma_i L \in (1, \overline{\gamma L}_1(\kappa)]$, we conjecture its tightness based on a two-dimensional worst-case example (Conjecture 3.6.1).

Remark 3.2.2 (Overlap of Theorem 3.2.3 and Theorem 3.2.4)**.** *Restricting it to constant stepsizes $\gamma L \leq \overline{\gamma L}_1(\kappa)$, Theorem 3.2.4 recovers the result from Theorem 3.2.3 corresponding to $\frac{-1+(1-\gamma\mu)^{-2k}}{\gamma\mu} - \frac{-1+(1-\gamma L)^{-2k}}{\gamma L} \leq 0$ for all $k \geq 1$ in (3.6), namely $P_N(\gamma L, \gamma\mu) = p(\gamma L, \gamma\mu)\, N$. The extension to variable stepsizes relies on the proof which, within the stepsize range from Theorem 3.2.4, only requires inequalities relating consecutive iterations.*

Figure 3.2 illustrates the behaviour of the leading term $\gamma L\, p(\gamma L, \gamma\mu)$, emphasizing the stepsize threshold $\overline{\gamma L}_1(\kappa)$ and additionally showing that optimal rates, i.e., with larger denominators, are achieved in the regime of stepsizes greater than 1.

Remark 3.2.3 (Recovering state of the art)**.** *Theorem 3.2.4 extends to arbitrary $\mu < 0$ the state-of-the-art rate for smooth not-necessarily convex functions ($\mu = -L$) from [2, Theorem 2], corresponding to stepsizes $\gamma_i L \in (0, \sqrt{3}]$:*

$$\frac{1}{2L} \min_{0 \leq i \leq N} \left\{ \|\nabla f(x_i)\|^2 \right\} \leq \frac{f(x_0) - f_*}{1 + \sum_{i=0}^{N-1} \gamma_i L \big(2 - \frac{\gamma_i L}{2} \max\{1, \gamma_i L\}\big)}. \tag{3.10}$$

The standard convergence rate (3.11)*, stated for instance in [89, §1.2.3], is derived for smooth functions using only the upper curvature L:*

$$\frac{1}{2L}\min_{0\le i\le N}\left\{\|\nabla f(x_i)\|^2\right\} \le \frac{f(x_0)-f_*}{1+\sum_{i=0}^{N-1}\gamma_i L(2-\gamma_i L)}, \quad \forall \gamma_i L \in (0,2). \tag{3.11}$$

Theorem 3.2.3 recovers this result in the absence of lower curvature information, i.e., for $\mu \searrow -\infty$ and effectively becomes Theorem 3.2.4 covering the full stepsize domain $(0,2)$ (as $\overline{\gamma L}_1(-\infty)=2$). The proof of this celebrated result lacks the utilization of quadratic lower bounds and thus remains generally non-tight. The convergence rate from Theorem 3.2.3 exhibits a smooth interpolation between the result in (3.11) *(when $\mu \searrow -\infty$), the rate for smooth not-necessarily convex functions ($\mu=-L$) from* (3.10) *and the rate on smooth convex functions ($\mu=0$).*

Corollary 3.2.2 (Exact performance bound for convex functions and variables stepsizes below 3/2)**.** *Let $f\in\mathcal{F}_{0,L}$ and consider N iterations of* (GD) *starting from x_0 with fixed stepsizes $\gamma_i L\in(0,\frac{3}{2}]$, $i=0,\dots,N-1$. Then the following bound holds:*

$$\frac{1}{2L}\|\nabla f(x_N)\|^2 \le \frac{f(x_0)-f_*}{1+2\sum_{i=0}^{N-1}\gamma_i L}.$$

Corollary 3.2.2 is obtained by taking $\mu=0$ in Theorem 3.2.4. Although not explicitly presented as such, Teboulle and Vaisbourd propose in [118] two sufficient decrease results (Lemma 1 and Lemma 5) yielding the same rate from Corollary 3.2.2 through telescoping summation (see also Remark 3.5.1).

3.2.5 Stepsize thresholds

Stepsize thresholds are given by the roots of (3.8). For N iterations, there are exactly N such thresholds, delimiting the regimes for $\gamma L>1$. These thresholds play a central role in the proofs for (strongly) convex and weakly convex functions in Theorems 3.2.1 to 3.2.3, as they partition the stepsize domain into $N+1$ distinct regimes, each corresponding to a different worst-case instance. We begin by introducing the following auxiliary expressions.

Definition 3.2.1 (T_k quantity determining the thresholds)**.** *Let k be a positive integer and define $E_k:\mathbb{R}\backslash\{0\}\to\mathbb{R}$:*

$$E_k(x) := \sum_{j=1}^{2k} x^{-j} = \begin{cases} \frac{-1+x^{-2k}}{1-x}, & x\neq 1; \\ 2k, & x=1. \end{cases} \tag{E_k}$$

Let $V \coloneqq \{(\ell,\kappa) :\ \kappa \in (-\infty, 1),\ \ell \in (1, \frac{2}{1+[\kappa]_+})\}$ *and define* $T_k : V \to \mathbb{R}$ *as*

$$T_k(\ell,\kappa) \coloneqq E_k\left(1-\kappa\ell\right) - E_k\left(1-\ell\right). \tag{T_k}$$

$T_k(\gamma L,\kappa)$ denotes exactly the individual terms summed up in (3.6) and canceled by the k-th stepsize threshold in (3.8). Moreover, $T_N(\gamma L,\kappa)$ is the difference of the arguments in the denominators of the rates for convex functions (3.1) and strongly convex functions (3.3), respectively.

Proposition 3.2.1 (Monotonicity of T_k)**.** *The function* $T_k(\cdot,\kappa)$ *is strictly increasing in the first argument.*

Proof. See Section 3.A.1. □

The sign of $T_k(\gamma L,\kappa)$ changes in the stepsize interval $\gamma L \in \left(1, \frac{2}{1+[\kappa]_+}\right)$, on which it therefore possesses a unique root.

Definition 3.2.2 (Stepsize thresholds)**.** *Let* $\overline{\gamma L}_\infty(\kappa) \coloneqq \frac{2}{1+[\kappa]_+}$. *The unique roots of* $T_k(\gamma L,\kappa)$ *in the interval* $(1, \overline{\gamma L}_\infty(\kappa))$ *are referred to as stepsize thresholds and denoted by* $\overline{\gamma L}_k(\kappa)$*, where* $\overline{\gamma L}_0(\kappa) \coloneqq 1$ *and:*

$$\left\{\overline{\gamma L}_k(\kappa)\right\} \coloneqq \left\{\gamma L \in (1, \overline{\gamma L}_\infty(\kappa)) \,|\, T_k(\gamma L,\kappa) = 0\right\} \quad \forall k = 1,2,\ldots. \tag{$\overline{\gamma L}_k$}$$

The definition of stepsize thresholds ($\overline{\gamma L}_k$) resembles exactly the condition (3.8) marking the transitions between the different regimes indexed by k in Theorem 3.2.3.

Proposition 3.2.2 (Properties of stepsize thresholds)**.** *The stepsize thresholds* $\overline{\gamma L}_k(\kappa)$ *satisfy the following properties:*

1. $\overline{\gamma L}_k(\kappa) < \overline{\gamma L}_{k+1}(\kappa)$*, for all integers* $k \geq 0$*;*
2. $\lim_{k\to\infty} \overline{\gamma L}_k(\kappa) = \overline{\gamma L}_\infty(\kappa)$*;*
3. For any $\gamma L \in \left[\overline{\gamma L}_k(\kappa), \overline{\gamma L}_{k+1}(\kappa)\right)$*:* $T_i(\gamma L,\kappa) \geq 0\ \forall i \leq k$*, and* $T_i(\gamma L,\kappa) < 0$ $\forall i \geq k+1$*.*

Proof. See Section 3.A.1. □

Figure 3.1 illustrates the dependence on κ of the first $N = 4$ stepsize thresholds. Complementary, Figure 3.3 details the rapid increase of the thresholds with the number of iterations, reaching: (i) $\frac{2}{1+\kappa}$ for strongly convex functions and (ii) 2 in the non-strongly convex case. The distance between two consecutive thresholds decreases when increasing their index. Figure 3.4 depicts $T_k(\gamma L,\kappa)$

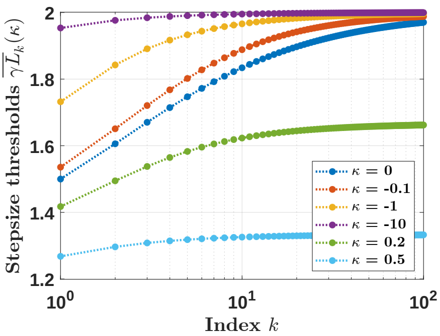

Figure 3.3: Stepsize thresholds $\overline{\gamma L}_k(\kappa)$ converge to $\frac{2}{1+\kappa}$ for strongly convex functions and to 2 for non-strongly convex functions.

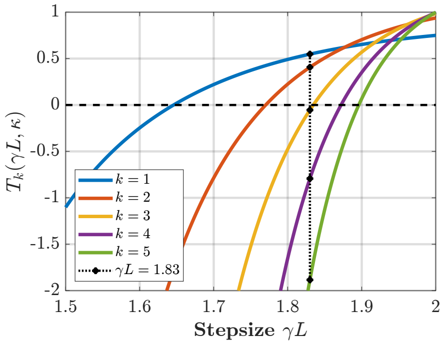

Figure 3.4: Dependence of $T_k(\gamma L, \kappa)$ on γL (for $\kappa = -0.5$): with $\gamma L = 1.83$, $T_k \geq 0$ for $k \in \{1, 2\}$. $\overline{\gamma L}_k(\kappa)$ are the intersections with the abscissa.

defined in (T_k) for $\kappa = -0.5$ and several indices k, offering insight on how to determine the stepsize thresholds. For example, with $\gamma L = 1.83$, $T_k(\gamma L, \kappa)$ is positive for $k \leq 2$ and negative for $k \geq 3$; thus, the interval corresponding to these conditions is $[\overline{\gamma L}_2, \overline{\gamma L}_3)$.

3.3 Stepsize schedules to improve worst-case guarantees

Building on our worst-case performance analysis, we provide several stepsize policies that minimize these upper bounds. Recommendations are provided for both constant stepsizes (Section 3.3.1), which depend on the iteration count, and variable stepsizes (Section 3.3.2); the latter improve upon the constant choices while also remaining independent of the horizon.

3.3.1 Optimal constant stepsizes

In this section, we derive the optimal constant stepsize that maximizes the worst-case performance established in Section 3.2. Proposition 3.3.1 and Proposition 3.3.2 provide this stepsize for convex and strongly convex functions, respectively. Proposition 3.3.3 derives the (asymptotically) optimal constant stepsize for weakly convex functions.

Proposition 3.3.1 (Optimal constant stepsize for convex functions)**.** *Let $f \in \mathcal{F}_{0,L}$ and consider N iterations of* (GD) *starting from x_0 with constant stepsizes. The worst-case rate for convex functions* (3.1) *is minimized by the choice $\gamma L = \overline{\gamma L}_N(0)$, implying the best worst-case bound:*

$$\frac{1}{2L}\|\nabla f(x_N)\|^2 \leq \frac{f(x_0) - f_*}{1 + 2N\,\overline{\gamma L}_N(0)}. \tag{3.12}$$

Proposition 3.3.2 (Optimal constant stepsize for strongly convex functions)**.** *Let $f \in \mathcal{F}_{\mu,L}$, with $\mu \in (0, L)$, and N iterations of* (GD) *starting from x_0 with constant stepsizes. The worst-case rates for strongly convex functions* (3.3) *is minimized by the choice $\gamma L = \overline{\gamma L}_N(\kappa)$, implying the best worst-case bound:*

$$\frac{1}{2L}\|\nabla f(x_N)\|^2 \leq \frac{f(x_0) - f_*}{(1 - \overline{\gamma L}_N(\kappa))^{-2N}}. \tag{3.13}$$

To the limit $N \to \infty$, $\overline{\gamma L}_N(\kappa)$ becomes $\frac{2}{1+\kappa}$, yielding a rate with $(\frac{1-\kappa}{1+\kappa})^{2N}$ behaviour.

The N-th stepsize threshold $\overline{\gamma L}_N(\kappa)$ is the best for (strongly) convex functions. When evaluated at $\overline{\gamma L}_N(\kappa)$, Definition 3.2.2 implies that the two arguments in the minimization defining the performance bounds of Theorem 3.2.1 and Theorem 3.2.2 coincide.

The same optimal constant stepsize $\overline{\gamma L}_N(\kappa)$ is minimizing the worst-case rate for the more standard performance criterion $\frac{f(x_N)-f_*}{\|x_0-x_*\|^2}$, since the denominators are the same (see Theorem 3.2.1 for convex functions and Theorem 3.2.2 for strongly convex functions). We refer the reader to the discussion following [44, Conjecture 3.1] for convex functions and [116, Sections 4.1.1-2] for strongly convex functions. These rates were recently proved in [68].

Proposition 3.3.3 (Optimal constant stepsize for weakly convex functions)**.** *Let $\kappa < 0$, $\bar{\kappa} := \frac{-9-5\sqrt{5}+\sqrt{190+90\sqrt{5}}}{4} \approx -0.1001$, and define $(\gamma L)^*(\kappa)$ as the unique solution belonging to the interval $[1, 2)$ of the equation:*

$$-\kappa(1+\kappa)(\gamma L)^3 + \left[3\kappa + (1+\kappa)^2\right](\gamma L)^2 - 4(1+\kappa)(\gamma L) + 4 = 0. \tag{3.14}$$

With respect to minimizing the tight performance bound for weakly convex functions from Theorem 3.2.3, $(\gamma L)^*(\kappa)$ is:

1. *the optimal constant stepsize for any N, when $\kappa \leq \bar{\kappa}$;*
2. *the asymptotically optimal constant stepsize as $N \to \infty$, i.e., when the algorithm runs for a sufficiently large number of iterations, if $\kappa \in (\bar{\kappa}, 0)$.*

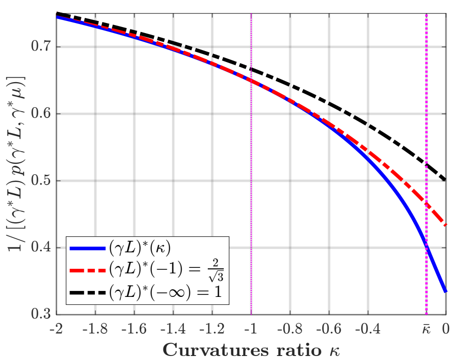

Figure 3.5: Comparison of benefit between (normalized) optimal stepsize recommendations: (i) $(\gamma L)^*(\kappa)$ from Proposition 3.3.3, (ii) $\gamma L = \frac{2}{\sqrt{3}}$ from [2] and (iii) $\gamma L = 1$ (classical).

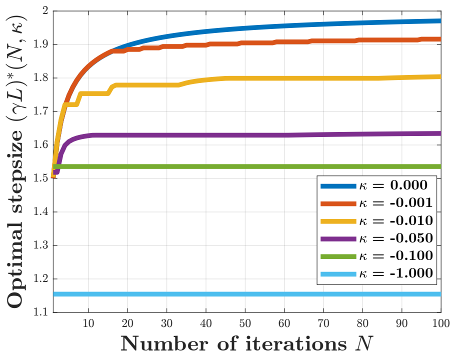

Figure 3.6: Optimal stepsize's (logarithmic) dependence on N: asymptotic convergence to $(\gamma L)^*(\kappa)$ from Proposition 3.3.3; closer to convex, a transient staircase-like behaviour emerges.

Proof. See Section 3.A.2. □

Unlike the optimal stepsize $\frac{2}{\sqrt{3}L}$ provided in [2, Theorem 3] for the usual smooth not-necessarily convex case $\mu = -L$, the recommendation in Proposition 3.3.3 leverages lower curvature information and maximizes the leading term $\gamma L\, p(\gamma L, \gamma\mu)$ in the rate's denominator from Theorem 3.2.3. As κ approaches $-\infty$, implying the absence of lower curvature information, the quantity $\gamma L(2 - \gamma L)$ is maximized, resulting in the celebrated optimal stepsize $\frac{1}{L}$. Similarly, substituting $\kappa = -1$ in equation (3.14) yields the optimal stepsize for smooth not-necessarily convex functions from [2].

Figure 3.1 depicts the continuous dependence of $(\gamma L)^*(\kappa)$ on the curvature ratio; above $\bar{\kappa} \approx -0.1$, the dashed line marks its asymptotic optimality. Figure 3.5 illustrates the maximum guaranteed improvement achievable by optimizing the leading term $p^{-1}(\gamma L, \gamma\mu)$ and compares: (i) the (asymptotically) optimal stepsize $(\gamma L)^*(\kappa)$ from Proposition 3.3.3, (ii) the optimal stepsize for smooth functions $(\gamma L)^*(-1) = \frac{2}{\sqrt{3}}$, and (iii) the celebrated optimal stepsize $(\gamma L)^*(-\infty) = 1$. As nonconvexity decreases ($\kappa \nearrow 0$), the optimal stepsize from Proposition 3.3.3 exhibits a greater improvement compared to stepsize $(\gamma L)^*(-1)$, which does not employ lower curvature information.

In contrast to the constant optimal stepsize recommendations for (strongly) convex functions, in the weakly convex case there exists an asymptotically optimal schedule independent of the number of iterations. A transient regime appears only for $-0.1 \lesssim \kappa < 0$, where the optimal constant stepsize also varies with N, i.e., $(\gamma L)^* = (\gamma L)^*(\kappa, N)$, increasing monotonically. Figure 3.6 shows this dependency for several curvature ratios, where the worst-case is numerically minimized for a given iteration budget. As $\kappa \nearrow 0$, the transient regime to reach the asymptotically optimal stepsize from Proposition 3.3.3 expands. During the initial iterations, a staircase-like behaviour is observed and the optimal stepsize is confined within a specific subdomain determined by the stepsize thresholds, i.e., $(\gamma L)^*(\kappa, N) \in [1, \overline{\gamma L}_N(\kappa))$. In the convex case, the staircase behaviour disappears, as the optimal constant stepsize $(\gamma L)^*(0, N)$ becomes $\overline{\gamma L}_N(0)$, lying at the intersection of two linear regimes (see Proposition 3.3.1).

3.3.2 Dynamic stepsizes

Stepsizes with $\gamma L > 1$ admit faster rates, motivating analysis in this regime. For (strongly) convex functions, variable stepsize policies can attain superior worst-case performance to constant ones by occasionally exceeding 2 [6, 5, 60, 139, 140], but they rely on intricate proofs, may depend on a fixed horizon, and their non-monotone nature makes them sensitive to curvature misestimation.

We introduce a monotonically increasing stepsize sequence inspired by Theorem 3.2.4 and the work of Teboulle-Vaisbourd [118], that avoids these issues. It applies to smooth strongly convex, convex, and weakly convex functions, is horizon-free, achieves better worst-case guarantees than the optimal constant stepsize, and admits significantly simpler convergence proofs. The construction is derived by telescoping a two-point descent inequality (Lemma 3.5.2), yielding a recurrence that cancels a key term and produces the largest admissible stepsize at each iteration.

Definition 3.3.1 (Dynamic stepsizes)**.** *Let $\kappa \in (-\infty, 1)$. We consider the sequence $\{s_k(\kappa)\}_{k=-1}^{\infty}$, with $s_{-1} = 0$, recursively defined as:*

$$s_{k+1}(\kappa) := \Big\{ s_+ \in \big(1, \tfrac{2}{1+[\kappa]_+}\big) : \frac{s_+[(2-s_+)(2-\kappa s_+)-1]}{2-s_+(1+\kappa)} + \frac{s_k(\kappa)}{2-s_k(\kappa)(1+\kappa)} = 0 \Big\}. \tag{3.15}$$

Proposition 3.3.4 (Properties of dynamic stepsizes)**.** *For any fixed $\kappa \in (-\infty, 1)$, the sequence $(s_k(\kappa))_{k=-1}^{\infty}$ is:*

(i) *uniquely defined;*
(ii) *monotonically increasing; and*
(iii) $\lim_{k\to\infty} s_k(\kappa) = \frac{2}{1+[\kappa]_+}$.

Moreover, for any $k \geq 1$ it holds that

(iv) $s_{k-1}(0) \geq 2 - \frac{1}{\frac{3}{2}k}$;
(v) $s_{k-1}(\kappa) \geq \frac{2}{1+\kappa} - (\frac{1-\kappa}{1+\kappa})^{2k}$, *if* $\kappa \in (0,1)$.

Proof. See Section 3.A.3. □

Note that $s_0 = \overline{\gamma L}_1(\kappa)$ is the largest value for which Theorem 3.2.4 holds. Obtaining a closed-form expression for s_k is difficult, as each step requires solving a quadratic equation (in the convex case) or cubic one (when $\kappa \neq 0$).

Theorem 3.3.1 (Performance bound with dynamic stepsizes)**.** *Let* $f \in \mathcal{F}_{\mu,L}$, *with* $\mu \in (-\infty, L)$, *and* $\kappa = \frac{\mu}{L}$. *Consider* N *iterations of* (GD) *starting from* x_0 *with stepsizes* $\gamma_i L$ *defined as follows:*

$$\gamma_i L \coloneqq \begin{cases} \min\{s_i(\kappa)\,,\,(\gamma L)^*(\kappa)\}, & \kappa < 0 \\ s_i(\kappa), & \kappa \geq 0 \end{cases}, \quad \forall\, i = 0, \ldots, N-1,$$

where $(\gamma L)^*(\kappa)$ *is defined in Proposition 3.3.3 and* $s_i(\kappa)$ *is given in Definition 3.3.1. Then the following bound holds:*

$$\frac{1}{2L} \min_{0 \leq i \leq N} \left\{\|\nabla f(x_i)\|^2\right\} \leq \frac{f(x_0) - f_*}{1 + \sum_{i=0}^{N-1} \frac{\gamma_i L(2-\gamma_i L)(2-\kappa\gamma_i L)}{2-\gamma_i L(1+\kappa)}}. \tag{3.16}$$

Proof. See Section 3.A.4. □

Based on numerical observations, we conjecture that the bound in (3.16) is tight. Theorem 3.3.1 extends Theorem 3.2.4 to larger stepsizes (i.e., $> \overline{\gamma L}_1(\kappa)$) and also applies to the strongly convex case ($\kappa > 0$). Theorem 3.3.1 is proved in Section 3.A.4, together with the resulting corollaries for convex, strongly convex and weakly convex functions, respectively (Corollaries 3.3.1 to 3.3.3).

Corollary 3.3.1 (Performance of dynamic stepsizes on convex functions)**.** *Let* $f \in \mathcal{F}_{0,L}$ *and consider* N *iterations of* (GD) *starting from* x_0 *with stepsizes* $\gamma_i L = s_i(0) = 1 + \frac{2}{1+\sqrt{9-4s_{i-1}(0)}}$ *for all* $i = 0, \ldots, N-1$. *Then the following bound holds:*

$$\frac{1}{2L}\|\nabla f(x_N)\|^2 \leq \frac{f(x_0) - f_*}{1 + \frac{s_{N-1}(0)}{2-s_{N-1}(0)}}. \tag{3.17}$$

According to Proposition 3.3.4(iv), the bound can be relaxed by $\frac{f(x_0)-f_*}{3N}$, which recovers the standard rate of $\mathcal{O}(\frac{1}{N})$.

Table 3.3: Comparison of complexity bounds (denominators) for convex functions between the policies: (i) celebrated stepsize $\gamma L = 1$; (ii) optimal constant stepsizes (Proposition 3.3.1) and (iii) dynamic schedule (from Corollary 3.3.1); higher is better. The last column displays the ratio of the these last two denominators.

N	Standard choice		Proposition 3.3.1		Corollary 3.3.1		
	$\gamma L = 1$	$1+2N$	$\overline{\gamma L}_N$	$1+2N\overline{\gamma L}_N$	s_{N-1}	$1+\frac{s_{N-1}}{2-s_{N-1}}$	Ratio (%)
1	1	3	1.500	4.000	1.500	4.000	100.000
2	1	5	1.606	7.423	1.732	7.464	100.549
5	1	11	1.747	18.471	1.893	18.619	100.806
10	1	21	1.834	37.681	1.947	37.933	100.670
20	1	41	1.897	76.885	1.974	77.235	100.456
30	1	61	1.924	116.426	1.983	116.826	100.343
40	1	81	1.939	156.106	1.987	156.535	100.275
50	1	101	1.949	195.859	1.990	196.310	100.230
100	1	201	1.971	395.109	1.995	395.612	100.127

Corollary 3.3.2 (Performance of dynamic stepsizes on strongly convex functions)**.** *Let $f \in \mathcal{F}_{\mu,L}$, with $\mu \in (0, L)$, and consider N iterations of* (GD) *starting from x_0 with stepsizes $\gamma_i L = s_i(\kappa)$ from Definition 3.3.1, with $i = 0, \ldots, N-1$. Then the following bound holds:*

$$\frac{1}{2L}\|\nabla f(x_N)\|^2 \leq \frac{f(x_0) - f_*}{1 + \frac{s_{N-1}(\kappa)}{2 - s_{N-1}(\kappa)(1+\kappa)}}. \tag{3.18}$$

As a consequence of Proposition 3.3.4(v), the gradient norm converges to zero at a rate of $\mathcal{O}((\frac{1-\kappa}{1+\kappa})^{2N})$.

Following the approach in [118, Table 1], we numerically compare our dynamic stepsize policy with the optimal constant stepsize dependent on the number of iterations. This comparison is presented in Table 3.3 (convex case) and Table 3.4 (strongly convex case). Such a numerical comparison is necessary to evaluate the worst-case performance, as our dynamic policy lacks a closed-form expression.

Corollary 3.3.3 (Performance of dynamic stepsizes on weakly convex functions)**.** *Let $f \in \mathcal{F}_{\mu,L}$, with $\mu \in (-\infty, 0)$, and consider N iterations of* (GD) *starting from x_0 with stepsizes $\gamma_i L = \min\{s_i(\kappa)\,,\,(\gamma L)^*(\kappa)\}$, with $i = 0, \ldots, N-1$, $s_i(\kappa)$*

Table 3.4: Comparison of complexity bounds (denominators) for strongly convex functions between the policies: (i) celebrated stepsize $\gamma L = \frac{2}{1+\kappa}$; (ii) optimal constant stepsizes (Proposition 3.3.2) and (iii) dynamic schedule (Corollary 3.3.2); higher is better. The last column displays the ratio of these last two denominators.

κ	N	Standard choice		Proposition 3.3.2		Corollary 3.3.2		
		$\frac{2}{1+\kappa}$	$\left(\frac{1-\kappa}{1+\kappa}\right)^{-2N}$	$\overline{\gamma L}_N$	$(1-\overline{\gamma L}_N)^{-2N}$	s_{N-1}	$1+\frac{s_{N-1}}{2-s_{N-1}(1+\kappa)}$	Ratio (%)
10^{-3}	1	1.998	1.004	1.500	4.006	1.500	4.006	100.000
	2	1.998	1.008	1.605	7.447	1.731	7.488	100.551
	5	1.998	1.020	1.746	18.633	1.892	18.784	100.813
	10	1.998	1.041	1.833	38.381	1.946	38.643	100.682
	20	1.998	1.083	1.896	79.879	1.973	80.257	100.473
	30	1.998	1.127	1.923	123.416	1.982	123.865	100.363
	40	1.998	1.174	1.938	168.871	1.986	169.373	100.297
	50	1.998	1.221	1.948	216.257	1.989	216.805	100.253
	70	1.998	1.492	1.960	317.040	1.992	317.670	100.199
10^{-4}	1	1.9998	1.000	1.500	4.001	1.500	4.001	100.000
	2	1.9998	1.001	1.606	7.426	1.732	7.467	100.550
	5	1.9998	1.002	1.747	18.487	1.893	18.636	100.807
	10	1.9998	1.004	1.834	37.750	1.947	38.004	100.671
	20	1.9998	1.008	1.897	77.178	1.974	77.531	100.457
	30	1.9998	1.012	1.924	117.101	1.983	117.505	100.345
	40	1.9998	1.016	1.939	157.323	1.987	157.759	100.277
	50	1.9998	1.020	1.949	197.780	1.990	198.240	100.232
	70	1.9998	1.028	1.961	279.309	1.993	279.801	100.176

defined in Definition 3.3.1 and $(\gamma L)^*(\kappa)$ given in Proposition 3.3.3. Then the following bound holds:

$$\frac{1}{2L}\min_{0\le i\le N}\left\{\|\nabla f(x_i)\|^2\right\} \le \frac{f(x_0)-f_*}{1+p_*N+\max\limits_{0\le k\le N}\left\{\frac{s_{k-1}(\kappa)}{2-s_{k-1}(\kappa)\,(1+\kappa)}-kp_*\right\}}, \quad (3.19)$$

where $p_* := p((\gamma L)^*, \kappa(\gamma L)^*)$ is the maximized coefficient $p(\gamma L, \gamma\mu)$ from (3.7).

In the weakly convex case, the sequence $\{s_k(\kappa)\}_{k=0}^{\infty}$ is truncated when exceeds the asymptotically optimal constant stepsize for nonconvex functions from Proposition 3.3.3. The numerical evidence in Table 3.5 supports that the rate from Corollary 3.3.3 outperforms both: (i) optimal constant stepsize schedule $(\gamma L)^*(N,\kappa)$ (obtained by numerically maximizing the denominator from Theorem 3.2.3) and (ii) asymptotically optimal constant stepsize $(\gamma L)^*(\kappa)$.

Table 3.5: Comparison of complexity bounds (denominators) for weakly convex functions (example $\kappa = -10^{-3}$) between the policies: (i) asymptotically constant optimal stepsize $(\gamma L)^*(\kappa)$ from Proposition 3.3.3, resulting in a denominator denoted by $P_N^A = P_N((\gamma L)^*(\kappa), \kappa(\gamma L)^*(\kappa))$; (ii) constant optimal stepsize depending on the number of iterations $(\gamma L)^*(N, \kappa)$ determined by numerically maximizing the denominator $P_N(\gamma L, \gamma\mu)$, denoted by $P_N^* = P_N((\gamma L)^*(N,\kappa), \kappa(\gamma L)^*(N,\kappa))$; and (iii) dynamic sequence from Corollary 3.3.3, leading to the denominator D_N. The transition between the increasing sequence and the constant schedule is marked by the horizontal line. The last column displays the ratio $D_N(\kappa)/P_N^*$.

N	Proposition 3.3.3		Proposition 3.3.3		Corollary 3.3.3		
	$(\gamma L)^*(\kappa)$	P_N^A	$(\gamma L)^*(N,\kappa)$	P_N^*	$\min\{s_{N-1}, (\gamma L)^*(\kappa)\}$	$D_N(\kappa)$	Ratio (%)
1	1.939	1.135	1.500	3.994	1.500	3.994	100.000
2	1.939	1.288	1.606	7.400	1.733	7.440	100.547
5	1.939	1.882	1.748	18.310	1.893	18.456	100.798
8	1.939	2.750	1.809	29.509	1.935	29.721	100.719
9	1.939	3.121	1.823	33.254	1.939	33.483	100.689
10	1.939	3.542	1.835	36.999	1.939	37.246	100.667
20	1.939	12.546	1.885	74.170	1.939	74.867	100.940
30	1.939	44.438	1.894	111.360	1.939	112.489	101.014
40	1.939	144.899	1.898	148.645	1.939	150.111	100.986
50	1.939	182.520	1.905	186.003	1.939	187.733	100.930
100	1.939	370.629	1.916	373.255	1.939	375.841	100.693

As κ approaches 0, $(\gamma L)^*(\kappa)$ converges to 2 and the result from Corollary 3.3.1 is retrieved.

3.4 Preliminary results for the proofs

The necessary and sufficient interpolation conditions for $\mathcal{F}_{\mu,L}$-functions given in Theorem 2.2.1 are used in conjunction with the (GD) iterations by substituting

$$x_j = x_i - \sum_{k=i}^{j-1} \gamma_k \nabla f(x_k),$$

as detailed in Definition 3.4.1. Since the proofs are limited to constant stepsizes, we set γ_k to γ.

Definition 3.4.1 (Distance-$(j-i)$ inequalities for gradient descent)**.** *Consider inequalities* $(Q_{[i,j]})$ *and* $(Q_{[j,i]})$ *connecting gradient descent iterations separated* *by $(j-i)$ steps:*

$$f(x_i) - f(x_j) \geq \gamma \Big\langle \nabla f(x_j), \sum_{k=i}^{j-1} \nabla f(x_k) \Big\rangle + \frac{1}{2L} \left\| \nabla f(x_i) - \nabla f(x_j) \right\|^2 + \frac{\mu}{2L(L-\mu)} \Big\| \nabla f(x_i) - \nabla f(x_j) - \gamma L \sum_{k=i}^{j-1} \nabla f(x_k) \Big\|^2; \tag{$Q_{[i,j]}$}$$

$$f(x_j) - f(x_i) \geq -\gamma \Big\langle \nabla f(x_i), \sum_{k=i}^{j-1} \nabla f(x_k) \Big\rangle + \frac{1}{2L} \left\| \nabla f(x_i) - \nabla f(x_j) \right\|^2 + \frac{\mu}{2L(L-\mu)} \Big\| \nabla f(x_i) - \nabla f(x_j) - \gamma L \sum_{k=i}^{j-1} \nabla f(x_k) \Big\|^2. \tag{$Q_{[j,i]}$}$$

We refer to inequalities ($Q_{[i,j]}$) and ($Q_{[j,i]}$) as distance-1 or distance-2 inequalities if the distances between their indices satisfy $|i-j| \leq 1$ or $|i-j| \leq 2$, respectively.

In contrast to traditional convergence study, e.g., [89], we separate the analysis concerning the optimal objective value f_* and solely focus on the function value gap up to step N, namely $f(x_0) - f(x_N)$. By making a slight modification, as detailed in the descent result from Lemma 3.4.1, we derive a convergence rate incorporating f_* (alternative proof of the lemma given in [89, §1.2.3]).

Lemma 3.4.1 (Sufficient decrease for gradient descent)**.** *Let $f \in \mathcal{F}_{\mu,L}$, with $\mu \in (-\infty, L)$. One iteration of (GD) with stepsize $\gamma L \in (0,2)$ decreases the function values and at each point x_i it holds:*

$$f(x_i) - f_* \geq \frac{1}{2L} \|\nabla f(x_i)\|^2. \tag{3.20}$$

Proof. From Lemma 2.2.1 with $x = x_{i+1}$ and $y = x_i$:

$$f(x_{i+1}) - f(x_i) + \gamma \|\nabla f(x_i)\|^2 \leq \tfrac{\gamma^2 L}{2} \|\nabla f(x_i)\|^2,$$

hence the function values decrease for $\gamma L \in (0,2)$ since

$$f(x_i) - f(x_{i+1}) \geq \tfrac{\gamma(2-\gamma L)}{2} \|\nabla f(x_i)\|^2,$$

and since it holds for any x_i and γ, we can write it for $\gamma = \frac{1}{L}$ and get:

$$f(x_i) - f_* \geq f(x_i) - f(x_{i+1}) \geq \tfrac{1}{2L} \|\nabla f(x_i)\|^2.$$

□

A convergence rate for the last gradient norm $\|\nabla f(x_N)\|$ is obtained for (strongly) convex functions, since the gradient norm is non-increasing, as shown for example in [118, Lemma 2]. Lemma 3.4.2 provides an alternative proof of this property.

Lemma 3.4.2 (Monotonicity of gradient norm for convex functions)**.** *Let $f \in \mathcal{F}_{0,L}$ and one iteration with stepsize $\gamma > 0$, connecting x_i and x_{i+1}. Then:*

$$\|\nabla f(x_i)\|^2 - \|\nabla f(x_{i+1})\|^2 \geq \tfrac{2-\gamma L}{\gamma L}\|\nabla f(x_i) - \nabla f(x_{i+1})\|^2. \tag{3.21}$$

Moreover, if $\gamma \leq \frac{2}{L}$, then the gradient norm is non-increasing.

Proof. By replacing the (GD) iteration $x_{i+1} = x_i - \gamma \nabla f(x_i)$ in the formula of cocoercivity of the gradient [89, Theorem 2.1.12]:

$$\begin{aligned}\tfrac{1}{L}\|\nabla f(x_i) - \nabla f(x_{i+1})\|^2 &\leq \langle \nabla f(x_{i+1}) - \nabla f(x_i),\, x_{i+1} - x_i\rangle \\ &= -\gamma\langle \nabla f(x_{i+1}) - \nabla f(x_i),\, \nabla f(x_i)\rangle.\end{aligned}$$

By amplifying with L and using

$$\langle \nabla f(x_{i+1}), \nabla f(x_i)\rangle = \tfrac{1}{2}\|\nabla f(x_i)\|^2 + \tfrac{1}{2}\|\nabla f(x_{i+1})\|^2 - \tfrac{1}{2}\|\nabla f(x_i) - \nabla f(x_{i+1})\|^2$$

we obtain

$$\begin{aligned}\tfrac{1}{2}\|\nabla f(x_i)\|^2 - \tfrac{1}{2}\|\nabla f(x_{i+1})\|^2 + \tfrac{1}{2}\|\nabla f(x_i) - \nabla f(x_{i+1})\|^2 \geq \\ \tfrac{1}{\gamma L}\|\nabla f(x_i) - \nabla f(x_{i+1})\|^2,\end{aligned}$$

which simplifies to (3.21). □

3.5 Proofs of performance bounds

In general, solving the associated PEP which provides (numerical) exact performance bounds in closed form is difficult. We propose a systematic procedure which eliminates the standard PEP numerical guessing by effectively canceling specific terms in a linear combination of a well-defined set of inequalities from Definition 3.4.1 to mimic the shape of the performance metric, as described in Remark 3.5.2.

The PEP framework facilitated the identification and validation of worst-case scenarios, streamlining the necessary inequalities in the proofs, the subsequent demonstrations being inspired by the numerical solutions (which are not unique), with an improved understanding. Additionally, the intricate range of sublinear regimes outlined in Theorem 3.2.3 could only be validated numerically, but could not be deduced from the PEP numerical analyses. For a comprehensive understanding of the PEP setup for gradient descent, we refer the interested reader to [116] or [2].

Similar proofs using distance-1 interpolation conditions to obtain tight convergence rates on different performance metrics are given in [114, Theorem 2.1, Section 3] where the proximal gradient descent is analysed, assuming strong convexity on the smooth function in the split.

While using "distance-2" inequalities is a valid method for proving tight bounds with stepsizes greater than $\overline{\gamma L}_1(\kappa)$, we mention that the only fundamental requirement is to use conditions connecting non-consecutive iterates. In fact, PEP experiments reveal several alternative strategies, employing: both distance-2 inequalities but with different multipliers, only distance-2 inequalities of the type $Q_{[i+2,i]}$, or inequalities relating the final iterate $Q_{[N,i]}$ [68]. Our specific approach is advantageous because of its unique cancellation technique, described in Remark 3.5.2, which offers a methodical derivation. However, beyond this procedural benefit, we have not yet developed a deeper intuition.

Proofs organization. The lemmas from Sections 3.5.1 to 3.5.3 are employed in the rest of the section to prove the convergence rates outlined in Section 3.2. Table 3.2 provides a summary of the specific lemmas utilized in the demonstrations and the corresponding worst-case function types. The proofs are organized based on the stepsize thresholds delimiting the different regimes, even for (strongly) convex functions where the rates are expressed using minima in the denominators.

Main idea: Target inequality. Inequalities ($Q_{[i,j]}$) and ($Q_{[j,i]}$) enter with nonnegative weights $\alpha_{[i,j]}$ and $\alpha_{[j,i]}$, respectively, to form a generalized inequality linking N iterations[4] such that there exist nonnegative weights $\sigma_i \geq 0$ satisfying:

$$\sum_{0\leq i<j\leq N} \left(\alpha_{[i,j]} Q_{[i,j]} + \alpha_{[j,i]} Q_{[j,i]}\right) : \; f(x_0) - f(x_N) \geq \sum_{i=0}^{N} \sigma_i \tfrac{1}{2L} \|\nabla f(x_i)\|^2. \quad (\clubsuit)$$

Proposition 3.5.1 (General procedure)**.** *Inequality* ($\clubsuit$) *implies:*

$$\frac{f(x_0) - f(x_N)}{\sum_{i=0}^{N} \sigma_i} \geq \frac{1}{2L} \min_{0\leq i\leq N} \left\{\|\nabla f(x_i)\|^2\right\},$$

[4] For linear combination of inequalities, we adopt the convention of writing the summed expression on the l.h.s., while stating the corresponding inequalities explicitly on the r.h.s..

with equality only if all gradient norms $\|\nabla f(x_i)\|^2$ *with* $\sigma_i > 0$ *are equal. Moreover,*

$$\frac{f(x_0) - f_*}{1 + \sum_{i=0}^{N} \sigma_i} \geq \frac{1}{2L} \min_{0\leq i\leq N} \{\|\nabla f(x_i)\|^2\}.$$

Proof. The first inequality results after taking minimum gradient norm in (♣). The adjustment in the denominator of the second inequality results from combining (3.20) with (♣) and taking again the minimum gradient norm. □

Following Proposition 3.5.1, all convergence proofs reduce to demonstrating an inequality of the form (♣) with positive weights σ_i. We utilize interpolation inequalities ($Q_{[i,j]}$) and ($Q_{[j,i]}$) with (i) $|j-i| = 1$ to prove Theorem 3.2.4 and (ii) $|j-i| \in \{1, 2\}$ to prove Theorem 3.2.3 and Theorem 3.2.2.

Notation. For readability reasons, we sometimes denote $f_i = f(x_i)$ and $g_i = \nabla f(x_i)$. Recall: $\kappa = \frac{\mu}{L}$ is the curvature ratio; $\gamma L \in (0,2)$ is the (normalized) stepsize and $\overline{\gamma L}_\infty(\kappa) = \frac{2}{1+[\kappa]_+}$.

3.5.1 Distance-1 lemmas for weakly convex functions

For $|j-i| = 1$, inequalities ($Q_{[i,j]}$) and ($Q_{[j,i]}$) simplify to:

$$f_i - f_{i+1} \geq \frac{\gamma\mu\,\gamma L - 2\gamma\mu + 1}{L-\mu}\frac{\|g_i\|^2}{2} + \frac{1}{L-\mu}\frac{\|g_{i+1}\|^2}{2} + \frac{\gamma L - 1}{L-\mu}\langle g_i, g_{i+1}\rangle; \quad (Q_{[i,i+1]})$$

$$f_{i+1} - f_i \geq \frac{\gamma\mu\gamma L - 2\gamma L + 1}{L-\mu}\frac{\|g_i\|^2}{2} + \frac{1}{L-\mu}\frac{\|g_{i+1}\|^2}{2} + \frac{\gamma\mu - 1}{L-\mu}\langle g_i, g_{i+1}\rangle. \quad (Q_{[i+1,i]})$$

Note that their l.h.s. only contains consecutive function values, whereas the r.h.s. is a combination of gradients, with symmetric coefficients in γL and $\gamma\mu$.

Lemma 3.5.1 (One-step descent with $\gamma \leq 1/L$)**.** *Let* $f \in \mathcal{F}_{\mu,L}$*, with* $\mu \in (-\infty, L)$*, and consider one iteration of* (GD) *with stepsize* $\gamma L \in (0,1]$*. Then the following inequality holds with equality only if* $\nabla f(x_i) = \nabla f(x_{i+1})$ *or* $\gamma L = 1$*:*

$$f(x_i) - f(x_{i+1}) \geq \frac{\gamma L\,\gamma\mu - 2\gamma\mu + \gamma L}{L-\mu}\frac{\|\nabla f(x_i)\|^2}{2} + \frac{\gamma L}{L-\mu}\frac{\|\nabla f(x_{i+1})\|^2}{2}. \quad \text{(N2SD)}$$

Furthermore, if $\mu \leq 0$*, then the gradient norms from* (N2SD) *are weighted by positive quantities.*

Proof. Inequality ($Q_{[i,i+1]}$) is reformulated by expressing the inner product in terms of gradient norms, i.e., $2\langle g_i, g_{i+1}\rangle = \|g_i\|^2 + \|g_{i+1}\|^2 - \|g_i - g_{i+1}\|^2$, as:

$$f_i - f_{i+1} \geq \frac{\gamma L\,\gamma\mu - 2\gamma\mu + \gamma L}{L-\mu}\frac{\|g_i\|^2}{2} + \frac{\gamma L}{L-\mu}\frac{\|g_{i+1}\|^2}{2} + \frac{1-\gamma L}{L-\mu}\frac{\|g_i - g_{i+1}\|^2}{2}. \tag{3.22}$$

Since $\gamma L \in (0,1]$, the mixed term can be neglected, resulting in inequality (N2SD), with equality only if $g_i = g_{i+1}$ or $\gamma L = 1$. The numerator of first term rewrites as $-\gamma\mu(1-\gamma L) + (\gamma L - \gamma\mu)$, hence it is positive when $\mu \leq 0$. □

Lemma 3.5.1 does not hold for stepsizes $\gamma L \in (1,2)$.

Lemma 3.5.2 (One-step descent with $\gamma \geq 1/L$)**.** *Let $f \in \mathcal{F}_{\mu,L}$, with, $\mu \in (-\infty, L)$, and consider one iteration of* (GD) *with stepsize $\gamma L \in [1, \overline{\gamma L}_\infty(\kappa))$. Then the following inequality holds:*

$$f(x_i) - f(x_{i+1}) \geq \frac{\gamma[(2-\gamma L)(2-\gamma\mu)-1]}{2-\gamma L-\gamma\mu}\frac{\|\nabla f(x_i)\|^2}{2} + \frac{\gamma}{2-\gamma L-\gamma\mu}\frac{\|\nabla f(x_{i+1})\|^2}{2}. \tag{N4SD}$$

When $\gamma L > 1$, equality holds only if:

$$\langle \nabla f(x_i), \nabla f(x_{i+1})\rangle = \frac{(1-\gamma\mu)(1-\gamma L)}{2-\gamma L-\gamma\mu}\|\nabla f(x_i)\|^2 + \frac{1}{2-\gamma L-\gamma\mu}\|\nabla f(x_{i+1})\|^2. \tag{3.23}$$

Furthermore, if $\gamma L \in [1, \overline{\gamma L}_1(\kappa)]$, then the gradient norms from (N4SD) *are weighted by nonnegative quantities.*

Proof. The mixed term in inequality (3.22) enters with a negative contribution for $\gamma L > 1$. To counterbalance it, the complementary interpolation inequality ($Q_{[i+1,i]}$) is used, written with the equivalent expression:

$$f_{i+1} - f_i \geq \frac{\gamma L\gamma\mu - 2\gamma L + \gamma\mu}{L-\mu}\frac{\|g_i\|^2}{2} + \frac{\gamma\mu}{L-\mu}\frac{\|g_{i+1}\|^2}{2} + \frac{1-\gamma\mu}{L-\mu}\frac{\|g_i - g_{i+1}\|^2}{2}.$$

Consider the sum of distance-1 inequalities $\bar{I}_i^1 := Q_{[i,i+1]} + Q_{[i+1,i]}$:

$$0 \geq \frac{2\gamma L\gamma\mu - \gamma L - \gamma\mu}{L-\mu}\frac{\|g_i\|^2}{2} + \frac{\gamma L + \gamma\mu}{L-\mu}\frac{\|g_{i+1}\|^2}{2} + \frac{2-\gamma L-\gamma\mu}{L-\mu}\frac{\|g_i - g_{i+1}\|^2}{2}. \tag{$\bar{I}_i^1$}$$

By weighting ($\bar{I}_i^1$) with $\beta(\gamma L, \gamma\mu) := \frac{\gamma L - 1}{2-\gamma L-\gamma\mu} \geq 0$ and summing it with (3.22), the coefficient of $\|g_i - g_{i+1}\|^2$ gets canceled and we obtain (N4SD). The equality condition (3.23) follows from the equality case of ($\bar{I}_i^1$). Nevertheless, the coefficient of $\|g_i\|^2$ is nonnegative for $\gamma L \in [1, \overline{\gamma L}_1(\kappa)]$, where $\overline{\gamma L}_1(\kappa)$ is the largest stepsize $\gamma L \in [1,2)$ for which: $(2-\gamma L)(2-\gamma\mu)-1 \geq 0$, implying nonnegativity of the weight of $\|\nabla f(x_i)\|^2$ in (N4SD). □

The threshold $\overline{\gamma L}_1(\kappa)$ has the explicit formula $\overline{\gamma L}_1(\kappa) = \frac{3}{1+\kappa+\sqrt{1-\kappa+\kappa^2}}$.

Remark 3.5.1 (Generalizations of sufficient decrease results)**.** *Inequalities* (N2SD) *and* (N4SD) *are extensions to weakly convex functions of the "double" sufficient decrease formula* (2SD) *[118, Lemma 1] and the "quadruple sufficient decrease" formula* (4SD) *[118, Lemma 5], respectively, obtained for convex functions ($\mu = 0$):*

$$f(x_i)-f(x_{i+1}) \geq \gamma\frac{\|\nabla f(x_i)\|^2}{2} + \gamma\frac{\|\nabla f(x_{i+1})\|^2}{2},\ \forall \gamma L \in (0,1]; \tag{2SD}$$

$$f(x_i)-f(x_{i+1}) \geq \frac{\gamma(3-2\gamma L)}{2-\gamma L}\frac{\|\nabla f(x_i)\|^2}{2} + \frac{\gamma}{2-\gamma L}\frac{\|\nabla f(x_{i+1})\|^2}{2},\ \forall \gamma L \in [1,2). \tag{4SD}$$

3.5.2 Distance-1 lemmas for strongly convex functions

Further on, the analysis focuses on constant stepsizes $\gamma L \in (1,2)$, interval only partially covered by Theorem 3.2.4.

Additional notation. For brevity, we sometimes use $\rho = \rho(\gamma L) := 1-\gamma L \in (-1,0)$, $\eta = \eta(\gamma\mu) := 1-\gamma\mu$. Then $T_k(\gamma L,\kappa) = E_k(1-\kappa\gamma L) - E_k(1-\gamma L)$ (see Definition 3.2.1), where $E_k(x) = \sum_{j=1}^{2k} x^{-j}$, can be equivalently written as $T_k(\rho,\eta) := T_k(\gamma L,\kappa)$, with $T_k(\rho,\eta) = E_k(\eta) - E_k(\rho)$. Recall $\overline{\gamma L}_\infty(\kappa) = \frac{2}{1+[\kappa]_+}$.

Definition 3.5.1 (Iteration threshold $\bar{N}$)**.** *Let* $\bar{N}(\gamma L,\kappa) := \max\{k \in \mathbb{N} \,|\, T_k(\gamma L,\kappa) \geq 0\}$.

Proposition 3.5.2 (Basic inequalities for η and ρ)**.** *Let* $\mu \in (-\infty, L)$ *and* $\gamma L \in \left(1, \overline{\gamma L}_\infty(\kappa)\right)$. *Then* $\eta > 0$, $\eta + \rho > 0$ *and* $\rho^{-2} > \eta^{-2}$.

Proof. Condition $\gamma L < \overline{\gamma L}_\infty(\kappa)$ implies $\gamma L < \frac{2}{1+[\kappa]_+}$, hence $0 < 2-\gamma L - \gamma L[\kappa]_+ \leq 2-\gamma L-\gamma L\kappa$, implying $\eta+\rho > 0$. Since $\rho = 1-\gamma L \in (-1,0)$, we have $\eta = 1-\gamma\mu > 0$. Condition $L > \mu$ implies $\eta > \rho$ and together with $\eta+\rho > 0$ we conclude that $\eta^2 > \rho^2$, therefore $\rho^{-2} > \eta^{-2}$. □

Lemma 3.5.3 (One-step descent in the linear regimes for $\gamma > 1/L$)**.** *Let $f \in \mathcal{F}_{\mu,L}$, with $\mu \in (-\infty, L)$, and consider one iteration of (GD) from x_i with stepsize $\gamma L \in (1,2)$. Then it holds:*

$$\begin{aligned}\frac{f(x_i) - f(x_{i+1})}{\gamma} \geq -E_i\left(\rho\right) \frac{\|\nabla f(x_i)\|^2}{2} + E_{i+1}\left(\rho\right) \frac{\|\nabla f(x_{i+1})\|^2}{2} + \\ + \frac{1-(\eta+\rho)E_{i+1}\left(\rho\right)}{\eta-\rho} \frac{\|\nabla f(x_{i+1}) - \rho \nabla f(x_i)\|^2}{2}.\end{aligned} \tag{3.24}$$

Furthermore, if $\mu \in (0, L)$ and $\gamma L \in [\frac{2}{1+\kappa}, 2)$, then $1-(\eta+\rho)E_i\left(\rho\right) > 0$.

Proof. Let $\beta_{i+1}(\rho) \coloneqq (-\rho) E_{i+1}\left(\rho\right) \;\; > \;\; 0$. Inequality (3.24) results after performing the simplifications in the linear combination of interpolation inequalities:

$$Q_{[i,i+1]} + \beta_{i+1}(\rho)\left(Q_{[i,i+1]} + Q_{[i+1,i]}\right),$$

where ($Q_{[i,j]}$) and ($Q_{[j,i]}$) are equivalently rewritten for the pair $(i, i+1)$ as:

$$(1+\beta_{i+1})\, Q_{[i,i+1]} \colon \; \frac{f_i - f_{i+1}}{\gamma} \geq (1+\rho)\frac{\|g_i\|^2}{2} + \frac{1}{\eta-\rho}\frac{\|g_{i+1} - \rho g_i\|^2}{2};$$

$$\beta_{i+1}\, Q_{[i+1,i]} \colon \; \frac{f_{i+1} - f_i}{\gamma} \geq -\frac{\|g_i\|^2}{2} - \frac{1}{\rho}\frac{\|g_{i+1}\|^2}{2} + \frac{\eta}{\rho}\frac{1}{\eta-\rho}\frac{\|g_{i+1} - \rho g_i\|^2}{2}.$$

With $\gamma L \in [\frac{2}{1+\kappa}, 2)$, we have $\eta+\rho \leq 0$ (see Proposition 3.5.2) and therefore $1-(\eta+\rho)E_{i+1}\left(\rho\right) > 0$. ☐

Lemma 3.5.4 (One-step descent in the linear regimes for $\gamma L < \overline{\gamma L}_\infty$)**.** *Let $f \in \mathcal{F}_{\mu,L}$, with $\mu \in (-\infty, L)$, and consider one iteration of (GD) from x_i with stepsize $\gamma L \in \left(0, \overline{\gamma L}_\infty(\kappa)\right)$. Then it holds:*

$$\begin{aligned}\frac{f(x_i) - f(x_{i+1})}{\gamma} \geq \; -E_i\left(\eta\right) \frac{\|\nabla f(x_i)\|^2}{2} + E_{i+1}\left(\eta\right) \frac{\|\nabla f(x_{i+1})\|^2}{2} + \\ + \frac{-1+(\eta+\rho)E_{i+1}\left(\eta\right)}{\eta-\rho} \frac{\|\nabla f(x_{i+1}) - \eta \nabla f(x_i)\|^2}{2}.\end{aligned} \tag{3.25}$$

Furthermore, if $\mu \in (0, L)$ and $\gamma L \in \left(0, \overline{\gamma L}_{i+1}(\kappa)\right)$, then $-1+(\eta+\rho)E_{i+1}\left(\eta\right) > 0$.

Proof. Firstly, note that $\eta = 1-\gamma\mu > 0$; the case $\gamma L \in \left(1, \frac{2}{1+[\kappa]_+}\right)$ is covered by Proposition 3.5.2, while for $\gamma L \in (0,1]$ we have $1-\gamma\mu > 1-\gamma L \geq 0$. Let:

$$\beta_{i+1}(\eta) \coloneqq -1 + \eta E_{i+1}\left(\eta\right) \; = \frac{1+\eta+\cdots+\eta^{2i}}{\eta^{2i+1}} > 0.$$

Inequality (3.25) results after performing the simplifications in the linear combination of interpolation inequalities:

$$Q_{[i,i+1]} + \beta_{i+1}(\eta)\,(Q_{[i,i+1]} + Q_{[i+1,i]}),$$

where ($Q_{[i,j]}$) and ($Q_{[j,i]}$) equivalently rewrite for the pair $(i, i+1)$ as:

$$(1+\beta_{i+1})\,Q_{[i,i+1]}\colon\ \frac{f_i - f_{i+1}}{\gamma} \geq \frac{\|g_i\|^2}{2} + \frac{1}{\eta}\frac{\|g_{i+1}\|^2}{2} + \frac{\frac{\rho}{\eta}}{\eta-\rho}\frac{\|g_{i+1}-\eta g_i\|^2}{2};$$

$$\beta_{i+1}\,Q_{[i+1,i]}\colon\ \frac{f_{i+1} - f_i}{\gamma} \geq -(1+\eta)\frac{\|g_i\|^2}{2} + \frac{1}{\eta-\rho}\frac{\|g_{i+1}-\eta g_i\|^2}{2}.$$

Condition $\gamma L < \overline{\gamma L}_{i+1}(\kappa)$ implies $T_{i+1}(\rho,\eta) < 0$, hence $i \geq \bar{N}(\gamma L, \kappa)$ (recall Definition 3.5.1: $T_{\bar{N}}(\rho,\eta) \geq 0$ and $T_{\bar{N}+1}(\rho,\eta) < 0$, corresponding to $\gamma L < \overline{\gamma L}_{\bar{N}+1}(\kappa)$). Then the coefficient of the mixed term in (3.25) is positive:

$$\frac{-1+(\eta+\rho)E_{i+1}(\eta)}{\eta-\rho} \geq \frac{-1+(\eta+\rho)E_{\bar{N}+1}(\eta)}{\eta-\rho} = \frac{T_{\bar{N}}(\rho,\eta) - \rho^2 T_{\bar{N}+1}(\rho,\eta)}{(\eta-\rho)^2} > 0.$$

The first inequality results from the monotonic increase with index i of $E_{i+1}(\eta)$, whereas the last one by the definition of $\bar{N}(\gamma L, \kappa)$. □

3.5.3 Distance-2 lemmas

In this section we derive descent lemmas involving distance-2 interpolation inequalities, needed to prove rates on the stepsize range $\gamma L \in (1, \overline{\gamma L}_\infty(\kappa))$.

Lemma 3.5.5 (Key distance-2 result)**.** *Let $f \in \mathcal{F}_{\mu,L}$, with $\mu \in (-\infty, L)$. Consider $k+1$ iterations of (GD) with constant stepsize $\gamma L \in \big[\overline{\gamma L}_k(\kappa), \overline{\gamma L}_\infty(\kappa)\big)$ starting from x_0, where $k \geq 1$. Then it holds:*

$$\frac{f(x_0)-f(x_k)}{\gamma} + \frac{(\eta+\rho)T_k(\rho,\eta)}{2(\eta-\rho)}\frac{f(x_k)-f(x_{k+1})}{\gamma} \geq \left(E_k\left(\eta\right) + E_k\left(\rho\right)\right)\frac{\|\nabla f(x_k)\|^2}{4} +$$

$$+ \frac{T_k(\rho,\eta)}{2(\eta-\rho)}\Big[\eta\rho\frac{\|\nabla f(x_k)\|^2}{2} + \langle \nabla f(x_k), \nabla f(x_{k+1})\rangle + \frac{\|\nabla f(x_{k+1})\|^2}{2}\Big]. \qquad \text{(D2)}$$

Proof. Consider two iterations connecting x_i, x_{i+1} and x_{i+2}, with $i \geq 0$. The interpolation inequalities ($Q_{[i,j]}$) and ($Q_{[j,i]}$) written for distance-1 ($Q_{[i,i+1]}$ and $Q_{[i+1,i]}$) and distance-2 ($Q_{[i,i+2]}$ and $Q_{[i+2,i]}$), respectively:

$$Q_{[i,i+1]}\colon\ \frac{f_i - f_{i+1}}{\gamma} \geq (1+\rho)\frac{\|g_i\|^2}{2} + \frac{1}{\eta-\rho}\frac{\|g_{i+1}-\rho g_i\|^2}{2};$$

$$Q_{[i+1,i]}\colon \frac{f_{i+1}-f_i}{\gamma} \geq -\frac{\|g_i\|^2}{2} - \frac{1}{\rho}\frac{\|g_{i+1}\|^2}{2} + \frac{\eta}{\rho}\frac{1}{\eta-\rho}\frac{\|g_{i+1}-\rho g_i\|^2}{2};$$

$$Q_{[i,i+2]}\colon \frac{f_i-f_{i+2}}{\gamma} \geq (1+\rho+\rho^2)\frac{\|g_i\|^2}{2} + \rho\frac{\|g_{i+1}\|^2}{2} + \frac{1}{\rho(\eta-\rho)}\cdot$$

$$\left[\rho(1-\eta)\frac{\|g_{i+1}-\rho g_i\|^2}{2} - (1-\rho)\frac{\|g_{i+2}-\rho g_{i+1}\|^2}{2} + \frac{\|g_{i+2}-\rho^2 g_i\|^2}{2}\right];$$

$$Q_{[i+2,i]}\colon \frac{f_{i+2}-f_i}{\gamma} \geq -\frac{\|g_i\|^2}{2} - \frac{1}{\rho}\frac{\|g_{i+1}\|^2}{2} - \frac{1+\rho}{\rho^2}\frac{\|g_{i+2}\|^2}{2} + \frac{1}{\rho(\eta-\rho)}\cdot$$

$$\left[\eta(1-\rho)\frac{\|g_{i+1}-\rho g_i\|^2}{2} - (1-\eta)\frac{\|g_{i+2}-\rho g_{i+1}\|^2}{2} + \frac{\eta}{\rho}\frac{\|g_{i+2}-\rho^2 g_i\|^2}{2}\right],$$

are summed up in a linear combination with *nonnegative* corresponding weights:

$$(3.26) \equiv \alpha_{[i,i+1]}Q_{[i,i+1]} + \alpha_{[i+1,i]}Q_{[i+1,i]} + \alpha_{[i,i+2]}Q_{[i,i+2]} + \alpha_{[i+2,i]}Q_{[i+2,i]},$$

where:

$$\alpha_{[i,i+1]} \coloneqq 1 - \rho E_{i+1}(\rho) + \frac{(1-\eta)T_i(\rho,\eta)}{2(\eta-\rho)} - \frac{\eta(1+\rho)T_{i+1}(\rho,\eta)}{2(\eta-\rho)};$$

$$\alpha_{[i+1,i]} \coloneqq -\rho E_{i+1}(\rho) + \frac{(1+\rho)T_i(\rho,\eta)}{2(\eta-\rho)} + \frac{\rho(1-\eta)T_{i+1}(\rho,\eta)}{2(\eta-\rho)};$$

$$\alpha_{[i,i+2]} \coloneqq \frac{\eta T_{i+1}(\rho,\eta)}{2(\eta-\rho)};$$

$$\alpha_{[i+2,i]} \coloneqq \frac{-\rho T_{i+1}(\rho,\eta)}{2(\eta-\rho)},$$

and (3.26) is obtained after performing the simplifications:

$$\begin{aligned}\frac{f_i-f_{i+1}}{\gamma} + (\eta+\rho)\Big[-\frac{T_i(\rho,\eta)}{2(\eta-\rho)}\frac{f_i-f_{i+1}}{\gamma} + \frac{T_{i+1}(\rho,\eta)}{2(\eta-\rho)}\frac{f_{i+1}-f_{i+2}}{\gamma}\Big] \geq \\ -\frac{T_i(\rho,\eta)}{2(\eta-\rho)}\Big[\eta\rho\frac{\|g_i\|^2}{2} + \langle g_i, g_{i+1}\rangle + \frac{\|g_{i+1}\|^2}{2}\Big] \\ +\frac{T_{i+1}(\rho,\eta)}{2(\eta-\rho)}\Big[\eta\rho\frac{\|g_{i+1}\|^2}{2} + \langle g_{i+1}, g_{i+2}\rangle + \frac{\|g_{i+2}\|^2}{2}\Big] \\ -(E_i(\eta)+E_i(\rho))\frac{\|g_i\|^2}{4} + (E_{i+1}(\eta)+E_{i+1}(\rho))\frac{\|g_{i+1}\|^2}{4}.\end{aligned} \tag{3.26}$$

Telescoping it for indices $i = 0, 1, \ldots, k-1$ yields inequality (D2).

It remains to show the nonnegativity of multipliers. Since the stepsize thresholds increase with index i, condition $\gamma L \geq \overline{\gamma L}_k(\kappa)$ ensures $T_{i+1}(\rho, \eta) \geq 0$, for all $i = 0, \ldots, k-1$. This directly implies the positivity of $\alpha_{[i,i+2]}$ and $\alpha_{[i+2,i]}$. The positivity of distance-1 multipliers is evidenced by their equivalent representations as nonnegative sums (recall $\rho \in (-1, 0)$, $\eta \in (0, \infty)$; $\eta + \rho > 0$ and $\rho^{-2} > \eta^{-2}$ from Proposition 3.5.2):

$$2\alpha_{[i,i+1]} = 1 + \frac{\eta + \eta^{-2i}}{1+\eta} + \frac{(1+\rho)(\eta+\rho-\rho\eta)}{\eta\rho^2}\frac{-1+\rho^{-2(i+1)}}{-1+\rho^{-2}} + \frac{(1+\rho)(1+\eta^2)}{\eta(\eta-\rho)}\sum_{j=0}^{i}\left(\rho^{-2j} - \eta^{-2j}\right) > 0;$$

$$2\alpha_{[i+1,i]} = \frac{-1+\rho^{-2i}}{1-\rho} + \frac{\eta(1+\rho)}{1-\rho}\frac{-1+\eta^{-2(i+1)}}{1-\eta} + \frac{-\rho\big[1-\rho+\eta(1+\rho)\big]}{1-\rho}\frac{\rho^{-2(i+1)} - \eta^{-2(i+1)}}{\eta-\rho} > 0.$$

□

Remark 3.5.2 (Derivation of key result Lemma 3.5.5)**.** *Inequality* (D2) *selectively involves gradient norms of indices $\{k, k+1\}$, function values of indices $\{0, k, k+1\}$ and the inner product $\langle \nabla f(x_k), \nabla f(x_{k+1})\rangle$. As a result, the linear combination of interpolation inequalities eliminates all other function values and gradients. Specifically, it removes: k gradient norms, k consecutive and k distance-2 inner products of gradients, $k-1$ function values, while setting the coefficient of f_0 to $\frac{1}{\gamma}$, totaling $4k$ constraints. Remarkably, the linear combination employs exactly $4k$ inequalities as well, highlighting why we needed to involve distance-2 interpolation inequalities. Consequently, deriving Lemma 3.5.5 (and subsequent convergence rates) entails solving a linear system, contrasting with the PEP numerical analysis, which involves solving an SDP.*

Subsequent lemmas involving distance-2 interpolation inequalities are derived by augmenting Lemma 3.5.5 with distance-1 interpolation inequalities.

Lemma 3.5.6 (Multistep descent with $\gamma L \in [\overline{\gamma L}_k, \overline{\gamma L}_\infty)$)**.** *Let $f \in \mathcal{F}_{\mu,L}$, with $\mu \in (-\infty, L)$, and consider $k+1 \geq 1$ iterations of* (GD) *with stepsize $\gamma L \in [\overline{\gamma L}_k(\kappa), \overline{\gamma L}_\infty(\kappa))$, satisfying $\gamma L > 1$, starting from x_0. Then the following* *inequality holds:*

$$\frac{f(x_0)-f(x_{k+1})}{\gamma}\geq\frac{-\eta^2\rho^2T_{k+1}(\rho,\eta)}{\eta^2-\rho^2}\frac{\|\nabla f(x_k)\|^2}{2}+\frac{T_k(\rho,\eta)+(\eta-\rho)}{\eta^2-\rho^2}\frac{\|\nabla f(x_{k+1})\|^2}{2}. \tag{GN4SD}$$

Furthermore, if $\mu\in(-\infty,0]$ *and* $\gamma L\in[\overline{\gamma L}_k(\kappa),\overline{\gamma L}_{k+1}(\kappa))$*, then the gradient norms from* (GN4SD) *are weighted by nonnegative quantities.*

Proof. For $k=0$, the result follows from Lemma 3.5.2 and we obtain exactly (N4SD). For $k\geq 1$, inequality (GN4SD) is derived by adjusting (D2) through the linear combination of interpolation inequalities $\alpha_{[k,k+1]}Q_{[k,k+1]}+\alpha_{[k+1,k]}Q_{[k+1,k]}$:

$$\begin{aligned}\Big[1-\frac{(\eta+\rho)\,T_k(\rho,\eta)}{2(\eta-\rho)}\Big]\frac{f_k-f_{k+1}}{\gamma}\geq\\ \Big(-E_k\left(\eta\right)-E_k\left(\rho\right)-\frac{\eta\rho T_k(\rho,\eta)}{\eta-\rho}-\frac{2\eta^2\rho^2T_{k+1}(\rho,\eta)}{\eta^2-\rho^2}\Big)\frac{\|g_k\|^2}{4}\\ +\Big(\frac{\eta-\rho+T_k(\rho,\eta)}{\eta^2-\rho^2}-\frac{T_k(\rho,\eta)}{2(\eta-\rho)}\Big)\frac{\|g_{k+1}\|^2}{2}-\frac{T_k(\rho,\eta)}{2(\eta-\rho)}\langle g_k,g_{k+1}\rangle,\end{aligned}$$

where the corresponding weights are defined as follows:

$$\alpha_{[k,k+1]}\coloneqq 1+\frac{-\rho}{\eta+\rho}+\frac{(1-\eta)(\eta+\rho)-2\rho}{2(\eta^2-\rho^2)}T_k(\rho,\eta);$$

$$\alpha_{[k+1,k]}\coloneqq\frac{-\rho}{\eta+\rho}+\frac{(1+\rho)(\eta+\rho)-2\rho}{2(\eta^2-\rho^2)}T_k(\rho,\eta).$$

Multiplier $\alpha_{[k+1,k]}$ is represented as a sum of positive quantities ($\rho<0$). Multiplier $\alpha_{[k,k+1]}$ is directly written as sum of positive quantities if $(1-\eta)(\eta+\rho)-2\rho>0$ (in particular, this covers the (strongly) convex case where $\eta\leq 1$). Otherwise, $\alpha_{[k,k+1]}$ is equivalently expressed as a sum of positive terms as follows, with $\eta>1$:

$$\begin{aligned}\alpha_{[k,k+1]}=\frac{\eta}{2(\eta+\rho)}+\frac{1}{2(\eta^2-\rho^2)}\Big[\frac{-\rho(\eta-\rho)(1+\eta^2)}{(\eta-1)(1-\rho)}+\\ -\big[(1-\eta)(\eta+\rho)-2\rho\big]\Big(\frac{\eta^{-2k}}{\eta-1}+\frac{\rho^{-2k}}{1-\rho}\Big)\Big]>0.\end{aligned}$$

□

Lemma 3.5.7 (Multistep descent in the linear regime with $\gamma L \geq \overline{\gamma L}_k$)**.** *Let $f \in \mathcal{F}_{\mu,L}$, with $\mu \in (-\infty, L)$, and consider $k+1 \geq 1$ iterations of* (GD) *with stepsize $\gamma L \in [\overline{\gamma L}_k(\kappa), \overline{\gamma L}_\infty(\kappa))$, satisfying $\gamma L > 1$, starting from x_0. Then the following inequality holds:*

$$\frac{f(x_0)-f(x_{k+1})}{\gamma} \geq E_{k+1}(\rho) \frac{\|\nabla f(x_{k+1})\|^2}{2} + \frac{\eta^2 T_{k+1}(\rho,\eta)}{(\eta-\rho)^2} \frac{\|\nabla f(x_{k+1}) - \rho \nabla f(x_k)\|^2}{2}. \tag{3.27}$$

Furthermore, if $\gamma L \in [\overline{\gamma L}_{k+1}(\kappa), \overline{\gamma L}_\infty(\kappa))$, then $T_{k+1}(\rho,\eta) \geq 0$.

Proof. For $k = 0$, the result follows from Lemma 3.5.3. For $k \geq 1$, inequality (3.27) is derived by adjusting (D2) with the linear combination of interpolation inequalities $\alpha_{[k,k+1]} Q_{[k,k+1]} + \alpha_{[k,k+1]} Q_{[k+1,k]}$:

$$\begin{aligned}
\Big[1 - \frac{(\eta+\rho)\, T_k(\rho,\eta)}{2(\eta-\rho)}\Big] \frac{f_k - f_{k+1}}{\gamma} &\geq \Big(-\frac{T_k(\rho,\eta)}{2(\eta-\rho)} - \frac{\eta^2 \rho T_{k+1}(\rho,\eta)}{(\eta-\rho)^2} \Big) \langle g_k, g_{k+1} \rangle + \\
&+ \Big[\frac{2\eta^2\rho^2 T_{k+1}(\rho,\eta)}{(\eta-\rho)^2} - E_k(\eta) - E_k(\rho) - \frac{\eta\rho T_k(\rho,\eta)}{\eta-\rho} \Big] \frac{\|g_k\|^2}{4} + \\
&+ \Big[\frac{\eta^2\, T_{k+1}(\rho,\eta)}{(\eta-\rho)^2} + E_{k+1}(\rho) - \frac{T_k(\rho,\eta)}{2(\eta-\rho)} \Big] \frac{\|g_{k+1}\|^2}{2},
\end{aligned}$$

where the corresponding weights are defined as:

$$\alpha_{[k,k+1]} := 1 + (-\rho) E_{k+1}(\rho) + \frac{1-\eta}{2(\eta-\rho)} T_k(\rho,\eta);$$

$$\alpha_{[k+1,k]} := (-\rho) E_{k+1}(\rho) + \frac{1+\rho}{2(\eta-\rho)} T_k(\rho,\eta).$$

Multiplier $\alpha_{[k+1,k]}$ is represented as a sum of positive quantities, while $\alpha_{[k,k+1]}$ is directly written as sum of positive terms if $\eta \leq 1$ (the (strongly) convex case). Otherwise, $\eta > 1$ and $\alpha_{[k,k+1]}$ is equivalently expressed as sum of positive terms as:

$$\alpha_{[k,k+1]} = \frac{1 + (-\rho)\rho^{-2(k+1)}}{2(1-\rho)} + \eta \frac{\eta\eta^{-2(k+1)} + (-\rho)\rho^{-2(k+1)}}{2(\eta-\rho)}.$$

□

Lemma 3.5.8 (Multistep descent in the linear regime with $\gamma L \in [\overline{\gamma L}_k, \overline{\gamma L}_\infty)$)**.** *Let $f \in \mathcal{F}_{\mu,L}$, with $\mu \in (-\infty, L)$, and consider $k+1 \geq 1$ iterations of* (GD) *with* *stepsize* $\gamma L \in [\overline{\gamma L}_k(\kappa), \overline{\gamma L}_\infty(\kappa))$*, satisfying* $\gamma L > 1$*, starting from* x_0*. Then the following inequality holds:*

$$\frac{f(x_0)-f(x_{k+1})}{\gamma} \geq E_{k+1}(\eta)\,\frac{\|\nabla f(x_{k+1})\|^2}{2} + \frac{-\rho^2 T_{k+1}(\rho,\eta)}{(\eta-\rho)^2}\,\frac{\|\nabla f(x_{k+1})-\eta \nabla f(x_k)\|^2}{2}. \tag{3.28}$$

Furthermore, if $\gamma L \in [\overline{\gamma L}_k(\kappa), \overline{\gamma L}_{k+1}(\kappa))$*, then* $T_{k+1}(\rho,\eta) < 0$*.*

Proof. Case $k = 0$ results from Lemma 3.5.4; further on, $k \geq 1$. Inequality (2.8) written for the pair $(i, j) = (k, k+1)$ reads:

$$0 \geq \frac{\eta\rho}{\eta-\rho}\|g_k\|^2 + \frac{1}{\eta-\rho}\|g_{k+1}\|^2 - \frac{\eta+\rho}{\eta-\rho}\langle g_k\,,\, g_{k+1}\rangle. \tag{3.29}$$

1. **Case** $\gamma L \in [\overline{\gamma L}_k(\kappa), \overline{\gamma L}_{k+1}(\kappa))$**.** By using $\eta > 0$, (3.29) equivalently rewrites as:

$$0 \geq -\eta^2\frac{\|g_k\|^2}{2} + \frac{\|g_{k+1}\|^2}{2} + \frac{\eta+\rho}{\eta-\rho}\frac{\|g_{k+1}-\eta g_k\|^2}{2}.$$

By multiplying it with $\frac{-\rho^2}{\eta^2-\rho^2}T_{k+1}(\rho,\eta)$ (positive because $T_{k+1}(\rho,\eta) < 0$) and adding it to (GN4SD), the coefficient of $\|g_k\|^2$ is canceled and we get:

$$\frac{f_0-f_{k+1}}{\gamma} \geq \frac{T_k(\rho,\eta)+(\eta-\rho)-\rho^2 T_{k+1}(\rho,\eta)}{\eta^2-\rho^2}\,\frac{\|g_{k+1}\|^2}{2} + \frac{-\rho^2 T_{k+1}(\rho,\eta)}{(\eta-\rho)^2}\,\frac{\|g_{k+1}-\eta g_k\|^2}{2},$$

which after simplifications is exactly (3.28).

2. **Case** $\gamma L \in \left[\overline{\gamma L}_{k+1}(\kappa), \overline{\gamma L}_\infty(\kappa)\right)$**.** After multiplying (3.29) with $\frac{-\rho\eta}{\eta-\rho} > 0$:

$$0 \geq \frac{\|g_{k+1}\|^2}{2} + \frac{-\eta^2}{(\eta-\rho)^2}\frac{\|g_{k+1}-\rho g_k\|^2}{2} + \frac{-\rho^2}{(\eta-\rho)^2}\frac{\|g_{k+1}-\eta g_k\|^2}{2}.$$

By scaling it with $T_{k+1}(\rho,\eta) > 0$ and summing it with (3.27), the coefficient of $\|g_{k+1}-\rho g_k\|^2$ is canceled and after simplifications we get (3.28).

□

3.5.4 Proof of performance bounds for weakly convex functions

Proof of Theorem 3.2.4. By telescoping the sufficient decrease inequalities (N2SD) and (N4SD) for indices $i = 1, \ldots, N$ and stepsizes $\gamma_0, \ldots, \gamma_{N-1}$, respectively:

$$f_0 - f_N \geq \frac{\|g_0\|^2}{2}\Big[\frac{\gamma_0 L\,\gamma_0\mu - 2\gamma_0\mu + \gamma_0 L}{L-\mu}\delta_0 + \frac{\gamma_0[(2-\gamma_0 L)(2-\gamma_0\mu)-1]}{2-\gamma_0 L-\gamma_0\mu}(1-\delta_0)\Big]$$

$$+ \sum_{i=1}^{N-1}\frac{\|g_i\|^2}{2}\Big[\frac{\gamma_{i-1}}{L-\mu}\delta_{i-1} + \frac{\gamma_{i-1}}{2-\gamma_{i-1}L-\gamma_{i-1}\mu}(1-\delta_{i-1}) +$$

$$+ \frac{\gamma_i L\,\gamma_i\mu - 2\gamma_i\mu + \gamma_i L}{L-\mu}\delta_i + \frac{\gamma_i[(2-\gamma_i L)(2-\gamma_i\mu)-1]}{2-\gamma_i L-\gamma_i\mu}(1-\delta_i)\Big]$$

$$+ \frac{\|g_N\|^2}{2}\Big[\frac{\gamma_{N-1}}{L-\mu}\delta_{N-1} + \frac{\gamma_{N-1}}{2-\gamma_{N-1}L-\gamma_{N-1}\mu}(1-\delta_{N-1})\Big], \tag{3.30}$$

where $\delta_i = 1$ if $\gamma_i \leq \frac{1}{L}$, and $\delta_i = 0$ otherwise. The telescoped inequality resembles a decrease inequality of type (♣) if all gradient norms have positive weights; a sufficient condition is $\frac{\gamma_i[(2-\gamma_i L)(2-\gamma_i\mu)-1]}{2-\gamma_i L-\gamma_i\mu} \geq 0$ for $i = 0,\dots,N-1$, which is satisfied by any choice of $\gamma_i L \in (0, \overline{\gamma L}_1(\kappa)]$ (see Lemma 3.5.2). By taking the minimum gradient norm and reordering:

$$f_0 - f_N \geq \min_{0\leq i\leq N}\Big\{\frac{\|g_i\|^2}{2L}\Big\}\Big[\sum_{i=0}^{N-1} L\Big(\frac{\gamma_i L}{L-\mu} + \frac{\gamma_i L\,\gamma_i\mu - 2\gamma_i\mu + \gamma_i L}{L-\mu}\Big)\delta_i +$$

$$+ \Big(\frac{\gamma_i L}{2-\gamma_i L-\gamma_i\mu} + \frac{\gamma_i L[(2-\gamma_i L)(2-\gamma_i\mu)-1]}{2-\gamma_i L-\gamma_i\mu}\Big)(1-\delta_i)\Big]$$

$$= \min_{0\leq i\leq N}\Big\{\frac{\|g_i\|^2}{2L}\Big\}\sum_{i=0}^{N-1}\gamma_i L\, p(\gamma_i L, \gamma_i\mu),$$

with $p(\gamma_i L, \gamma_i\mu)$ defined in (3.7). □

Proof of Theorem 3.2.3. The interval $(0, \overline{\gamma L}_1(\kappa)]$ is addressed by the proof of Theorem 3.2.4, so it remains to prove the rates for constant stepsizes in the range

$$[\overline{\gamma L}_1(\kappa), 2) = \bigcup_{k=1}^{N-1}[\overline{\gamma L}_k(\kappa), \overline{\gamma L}_{k+1}(\kappa)) \cup [\overline{\gamma L}_N(\kappa), 2).$$

1. **Case $\gamma L \in [\overline{\gamma L}_k(\kappa), \overline{\gamma L}_{k+1}(\kappa))$, with some $k \in \{1,\dots,N-1\}$.** To inequality (GN4SD) we append $(N-k-1)$ inequalities (N4SD), written for indices $i =$

$k+1,\ldots,N-1$. Then we get (using the usual notation with γ, L, μ):

$$\begin{aligned}\frac{f_0-f_N}{\gamma}\geq&\frac{-(1-\gamma L)^2(1-\gamma\mu)^2T_{k+1}(\gamma L,\kappa)}{(1-\gamma\mu)^2-(1-\gamma L)^2}\frac{\|g_k\|^2}{2}+\frac{T_k(\gamma L,\kappa)+\gamma L-\gamma\mu}{(1-\gamma\mu)^2-(1-\gamma L)^2}\frac{\|g_{k+1}\|^2}{2}\\&+\frac{(2-\gamma L)(2-\gamma\mu)-1}{2-\gamma L-\gamma\mu}\frac{\|g_{k+1}\|^2}{2}+\frac{(2-\gamma L)(2-\gamma\mu)}{2-\gamma L-\gamma\mu}\sum_{i=k+2}^{N-1}\frac{\|g_i\|^2}{2}\\&+\frac{1}{2-\gamma L-\gamma\mu}\frac{\|g_N\|^2}{2},\end{aligned}$$

simplifying to

$$\begin{aligned}\frac{f_0-f_N}{\gamma}\geq&\frac{-(1-\gamma L)^2(1-\gamma\mu)^2T_{k+1}(\gamma L,\kappa)}{(1-\gamma\mu)^2-(1-\gamma L)^2}\frac{\|g_k\|^2}{2}+\frac{T_k(\gamma L,\kappa)}{(1-\gamma\mu)^2-(1-\gamma L)^2}\frac{\|g_{k+1}\|^2}{2}\\&+\frac{(2-\gamma L)(2-\gamma\mu)}{2-\gamma L-\gamma\mu}\frac{\|g_{k+1}\|^2}{2}+\frac{(2-\gamma L)(2-\gamma\mu)}{2-\gamma L-\gamma\mu}\sum_{i=k+2}^{N-1}\frac{\|g_i\|^2}{2}\\&+\frac{1}{2-\gamma L-\gamma\mu}\frac{\|g_N\|^2}{2},\end{aligned}$$

which is of type ($\clubsuit$) since all gradient norms have positive coefficients because $T_k(\gamma L,\kappa)\geq 0$ and $T_{k+1}(\gamma L,\kappa)<0$. By taking the minimum gradient norm:

$$\min_{k\leq i\leq N}\left\{\frac{\|g_i\|^2}{2L}\right\}\leq\frac{f_0-f_N}{\gamma L\,P_N(\gamma L,\gamma\mu)},$$

where

$$\begin{aligned}P_N(\gamma L,\gamma\mu)=&\frac{-(1-\gamma L)^2(1-\gamma\mu)^2T_{k+1}(\gamma L,\kappa)+T_k(\gamma L,\kappa)+\gamma L-\gamma\mu}{(1-\gamma\mu)^2-(1-\gamma L)^2}+\\&\frac{(2-\gamma L)(2-\gamma\mu)}{2-\gamma L-\gamma\mu}(N-k-1).\end{aligned}$$

Using algebraic manipulations one can check the following equivalent expressions for $P_N(\gamma L,\gamma\mu)$ discussed in Section 3.2.3:

$$\begin{aligned}P_N(\gamma L,\gamma\mu)&=\frac{\frac{-1+(1-\gamma L)^{-2k}}{\gamma L[1-(1-\gamma L)^2]}-\frac{-1+(1-\gamma\mu)^{-2k}}{\gamma\mu[1-(1-\gamma\mu)^2]}}{\frac{1}{1-(1-\gamma L)^2}-\frac{1}{1-(1-\gamma\mu)^2}}+\frac{(2-\gamma L)(2-\gamma\mu)}{2-\gamma L-\gamma\mu}(N-k)\\&=\frac{(2-\gamma L)(2-\gamma\mu)}{2-\gamma L-\gamma\mu}\left[N-\frac{-\gamma\mu\gamma L}{\gamma L-\gamma\mu}\left[\frac{\frac{-1+(1-\gamma\mu)^{-2k}}{1-(1-\gamma\mu)^2}-k}{\gamma\mu}-\frac{\frac{-1+(1-\gamma L)^{-2k}}{1-(1-\gamma L)^2}-k}{\gamma L}\right]\right].\end{aligned}$$

The expression from (3.6) is obtained from the identity:

$$\frac{\frac{-1+(1-\gamma\mu)^{-2k}}{1-(1-\gamma\mu)^2}-k}{\gamma\mu}-\frac{\frac{-1+(1-\gamma L)^{-2k}}{1-(1-\gamma L)^2}-k}{\gamma L}=$$

$$\sum_{i=0}^{k}\left[\frac{-1+(1-\gamma\mu)^{-2i}}{\gamma\mu}-\frac{-1+(1-\gamma\mu)^{-2i}}{\gamma\mu}\right],$$

where in the r.h.s. are summed up quantities $T_i(\gamma L,\kappa)$ with $i=\{0,1,\dots,k\}$, i.e., for all indices for which they are positive. Therefore, the r.h.s. can be rewritten using the positive quantities over N terms such as in formula (3.6).
2. **Case** $\gamma L\in[\overline{\gamma L}_N(\kappa),2)$**.** From Lemma 3.5.7 with $k=N-1$ it holds:

$$\frac{f_0-f_N}{\gamma}\geq E_N\,(1-\gamma L)\,\frac{\|g_N\|^2}{2}+\frac{(1-\gamma\mu)^2\,T_N(\gamma L,\kappa)}{(\gamma L-\gamma\mu)^2}\,\frac{\|g_N-(1-\gamma L)g_{N-1}\|^2}{2},$$

where the mixed term can be neglected since $T_N(\gamma L,\kappa)\geq 0$ and obtain:

$$f_0-f_N\geq\left[-1+(1-\gamma L)^{-2N}\right]\frac{\|g_N\|^2}{2L},$$

which is of type (♣) and implies the linear regime from Theorem 3.2.3.

□

3.5.5 Proofs of performance bounds for strongly convex functions

If $\mu>0$, then $\overline{\gamma L}_\infty(\kappa)=\frac{2}{1+\kappa}<2$.

Proof of Theorem 3.2.2. We split the proof over stepsize intervals:

$$(0,2)=\left(0,\overline{\gamma L}_1(\kappa)\right]\cup\left(\bigcup_{k=1}^{N-1}\left[\overline{\gamma L}_k(\kappa),\overline{\gamma L}_{k+1}(\kappa)\right)\right)\cup\left[\overline{\gamma L}_N(\kappa),\tfrac{2}{1+\kappa}\right)\cup\left[\tfrac{2}{1+\kappa},2\right).$$

1. **Case** $\gamma L\in\left(0,\overline{\gamma L}_1(\kappa)\right]$. By telescoping (3.25) for indices $i=0,\dots,N-1$:

$$\frac{f_0-f_N}{\gamma}\geq E_N\,(\eta)\,\frac{\|g_N\|^2}{2}+\sum_{i=0}^{N-1}\frac{-1+(\eta+\rho)E_{i+1}\,(\eta)}{\eta-\rho}\,\frac{\|g_{i+1}-\eta g_i\|^2}{2},$$

where all mixed terms have nonnegative coefficients because $\gamma L\in\left(0,\overline{\gamma L}_1(\kappa)\right]$ (see Lemma 3.5.4). Hence they can be neglected to obtain the lower bound of type (♣):

$$\frac{f_0-f_N}{\gamma}\geq E_N\,(\eta)\,\frac{\|g_N\|^2}{2}. \tag{3.31}$$

2. **Case** $\gamma L \in \left[\overline{\gamma L}_k(\kappa), \overline{\gamma L}_{k+1}(\kappa)\right)$, with some $k \in \{1, \dots, N-1\}$. By telescoping (3.25) for indices $i = k+1, \dots, N-1$:

$$\frac{f_{k+1} - f_N}{\gamma} \geq -E_{k+1}(\eta) \frac{\|g_{k+1}\|^2}{2} + E_N(\eta) \frac{\|g_N\|^2}{2} + \sum_{i=k+1}^{N-1} \frac{-1+(\eta+\rho)E_{i+1}(\eta)}{\eta-\rho} \frac{\|g_{i+1}-\eta g_i\|^2}{2},$$

with nonnegative weights of the mixed terms since $\gamma L \in \left(0, \overline{\gamma L}_{k+1}(\kappa)\right)$ (see Lemma 3.5.4). After appending it to (3.28) from Lemma 3.5.8, we obtain:

$$\frac{f_0 - f_N}{\gamma} \geq E_N(\eta) \frac{\|g_N\|^2}{2} + \frac{-\rho^2 T_{k+1}(\rho, \eta)}{(\eta-\rho)^2} \frac{\|g_{k+1} - \eta g_k\|^2}{2} + \sum_{i=k+1}^{N-1} \frac{-1 + (\eta+\rho)E_{i+1}(\eta)}{\eta - \rho} \frac{\|g_{i+1} - \eta g_i\|^2}{2}.$$

Since $T_{k+1}(\rho, \eta) < 0$, all mixed terms can be neglected to get a (♣)-type inequality:

$$\frac{f_0 - f_N}{\gamma} \geq E_N(\eta) \frac{\|g_N\|^2}{2}. \tag{3.32}$$

3. **Case** $\gamma L \in \left[\overline{\gamma L}_N(\kappa), \frac{2}{1+\kappa}\right)$**.** From Lemma 3.5.7 with $k = N-1$ it holds:

$$\frac{f_0 - f_N}{\gamma} \geq E_N(\rho) \frac{\|g_N\|^2}{2} + \frac{\eta^2 T_N(\rho, \eta)}{(\eta-\rho)^2} \frac{\|g_N - \rho g_{N-1}\|^2}{2}.$$

Since $T_N(\gamma L, \kappa) \geq 0$, by neglecting the mixed term we obtain a (♣)-type inequality:

$$\frac{f_0 - f_N}{\gamma} \geq E_N(\rho) \frac{\|g_N\|^2}{2}. \tag{3.33}$$

4. **Case** $\gamma L \in \left[\frac{2}{1+\kappa}, 2\right)$**.** By telescoping Lemma 3.5.3 for $i = 0, 1, \dots, N-1$ it holds:

$$\frac{f_0 - f_N}{\gamma} \geq E_N(\rho) \frac{\|g_i\|^2}{2} + \sum_{i=0}^{N-1} \frac{1 - (\eta+\rho)E_{i+1}(\rho)}{\eta - \rho} \frac{\|g_{i+1} - \rho g_i\|^2}{2},$$

where the mixed terms can be neglected and then the inequality reduces to (3.33).

Since the transition between the two linear regimes is given by the sign of $T_N(\rho,\eta) = E_N(\eta) - E_N(\rho)$, we can glue together expressions (3.31), (3.32) and (3.33):

$$\frac{f_0 - f_N}{\gamma} \geq \min\{E_N(\eta)\ ,\ E_N(\rho)\}\frac{\|g_N\|^2}{2},$$

equivalent to:

$$f_0 - f_N \geq \min\left\{\frac{-1+(1-\gamma\mu)^{-2N}}{\mu}\ ,\ \frac{-1+(1-\gamma L)^{-2N}}{L}\right\}\frac{\|g_N\|^2}{2},$$

which implies the bound (3.3) and the complementary one:

$$\frac{1}{2L}\|\nabla f(x_N)\|^2 \leq \frac{f(x_0)-f(x_N)}{\min\left\{\frac{L}{\mu}\left[-1+(1-\gamma\mu)^{-2N}\right]\ ,\ -1+(1-\gamma L)^{-2N}\right\}}.$$

□

3.5.6 Results for convex functions

Corollary 3.5.1 (Multistep descent for convex functions)**.** *Let $f \in \mathcal{F}_{0,L}$. Consider $k+1 \geq 1$ iterations of* (GD) *with stepsize $\gamma L \in [\overline{\gamma L}_k(\kappa), \overline{\gamma L}_{k+1}(\kappa))$, satisfying $\gamma L > 1$, starting from x_0. Then the following inequality holds, with positive coefficients of gradient norms:*

$$f(x_0) - f(x_{k+1}) \geq \left[2(k+1)\frac{-(1-\gamma L)^2}{2-\gamma L} + \sum_{i=0}^{k}(1-\gamma L)^{-2i}\right]\frac{\|\nabla f(x_k)\|^2}{2L} + \\ + \left[2(k+1)\frac{1}{2-\gamma L} - \sum_{i=0}^{k}(1-\gamma L)^{-2i}\right]\frac{\|\nabla f(x_{k+1})\|^2}{2L}. \tag{G4SD}$$

Moreover, (G4SD) *is valid on the extended stepsize range $\gamma L \in (1,2)$.*

Proof. It results by taking $\eta = 1$ in (GN4SD). □

(G4SD) is a multistep generalization of the descent lemma (4SD) from [118].

Proof of Theorem 3.2.1. Letting $\mu \nearrow 0$, all terms summed in (3.6) become zero, simplifying to $P_N(\gamma L, 0) = 2N$, so that all sublinear regimes corresponding to stepsizes up to $\overline{\gamma L}_N(0)$ share the same expression. For stepsizes $\gamma L \geq$

$\overline{\gamma L}_N(0)$, it holds the linear regime following the expression in the denominator $E_N(1-\gamma L) = \frac{-1+(1-\gamma L)^{-2N}}{\gamma L}$. Additionally, the thresholds are determined by the roots of:

$$T_k(\gamma L, 0) = 2k - \frac{-1+(1-\gamma L)^{-2k}}{\gamma L}, \ \forall k = 1, \dots, N,$$

hence the rate can be expressed as a minimum in the denominator of (3.1). Same bound is recovered to the limit $\mu \searrow 0$ in strongly convex functions' rate (3.3). □

3.6 Tightness of performance bounds

The tightness of convergence rates is typically assessed by identifying a worst-case function matching the upper bound when iterating the algorithm. This function can be determined either analytically or by providing an interpolable set of triplets. Table 3.2 summarizes the various demonstrations of tightness based on different stepsizes ranges.

Remark 3.6.1 (Checking tightness)**.** *The exactness of all proved convergence rates can be confirmed numerically by solving the associated performance estimation problems (PEPs), for example by using the specialized software packages PESTO [115] (for MATLAB) or PEPit [56] (for Python). In this section we propose analytical proofs.*

For stepsizes in the range $\gamma L \in [\overline{\gamma L}_N(\kappa), 2)$, a single performance bound applies across all function classes (weakly convex, convex and strongly convex). As shown in Proposition 3.6.1, this unified bound is proven to be tight using the same worst-case function $L\frac{\|x\|^2}{2}$, which, for simplicity, can be reduced to one dimension.

Proposition 3.6.1 (Tightness for $\gamma L \in [\overline{\gamma L}_N(\kappa), 2)$)**.** *Let $L > 0$, N be a positive integer, $\gamma L \in [\overline{\gamma L}_N(\kappa), 2)$ and $f_L(x) := L\frac{\|x\|^2}{2}$. Then after applying N iterations of* (GD) *on function $f = f_L$ with constant stepsize γL, starting from x_0, it holds:*

$$\tfrac{1}{2L} \min_{0 \le i \le N} \{\|\nabla f(x_i)\|^2\} = \tfrac{1}{2L}\|\nabla f(x_N)\|^2 = \frac{f(x_0) - f_*}{(1-\gamma L)^{-2N}}.$$

Proof. One can directly check that f_L satisfies the necessary conditions for equality in Lemma 3.5.7: $g_N = (1-\gamma L) g_{N-1}$ and $f_0 - f_N = [-1 + (1-\gamma L)^{-2N}]\frac{\|g_N\|^2}{2L}$. □

For the remainder of the stepsize interval, $\gamma L \in (0, \overline{\gamma L}_N(\kappa))$, the proofs are organized as follows: Section 3.6.1 covers constant stepsizes for convex and strongly convex functions; Section 3.6.3 addresses variable stepsizes below $\overline{\gamma L}_1(\kappa)$ for convex and weakly convex functions; Section 3.6.4 focuses on the specific range of constant stepsizes $(1, \overline{\gamma L}_N(\kappa))$ for weakly convex functions.

3.6.1 Tightness for convex and strongly convex functions

Proposition 3.6.2 (Tightness of Theorem 3.2.2 (strongly convex))**.** *Let $L > 0$, $\mu \in (0, L)$, $\gamma L \in (0, 2)$ and N be a positive integer. Then there exists x_0 and $f \in \mathcal{F}_{\mu,L}$ such that after applying N iterations of (GD) on function f with constant stepsize $\gamma L \in (0, 2)$, starting from x_0, it holds:*

$$\tfrac{1}{2L}\|\nabla f(x_N)\|^2 = \frac{f(x_0) - f_*}{1 + \gamma L \min\left\{\frac{-1+(1-\gamma\mu)^{-2N}}{\gamma\mu}, \frac{-1+(1-\gamma L)^{-2N}}{\gamma L}\right\}}.$$

Proof. We define the following worst-case function depending on the stepsizes γL: if $\gamma L \in [\overline{\gamma L}_N(\kappa), 2)$, then it is given by the function f_L (defined in Proposition 3.6.1); if $\gamma L \in (0, \overline{\gamma L}_N(\kappa))$, then it is the function φ given below. Let $\Delta > 0$ and define the parameter $\tau := \sqrt{\frac{2}{L}\frac{\Delta}{1+\frac{L}{\mu}[-1+(1-\gamma\mu)^{-2N}]}}$. The function $\varphi\colon \mathbb{R} \to \mathbb{R}$ is given by:

$$\varphi(x) := \begin{cases} \mu\frac{x^2}{2} + (L-\mu)\tau|x| - (L-\mu)\frac{\tau^2}{2}, & \text{if } |x| \geq \tau; \\ L\frac{x^2}{2}, & \text{if } |x| < \tau. \end{cases}$$

Note that $\varphi \in \mathcal{F}_{\mu,L}$, $x_* = \arg\min \varphi(x) = 0$, $\varphi(x_*) = 0$ and

$$\nabla\varphi(x) := \begin{cases} \mu x + (L-\mu)\tau\,\mathrm{sgn}(x), & \text{if } |x| \geq \tau; \\ Lx, & \text{if } |x| < \tau. \end{cases}$$

One can check that initializing the iterates at $x_0{:=}\tau[1{+}\frac{L}{\mu}\left(-1{+}(1{-}\gamma\mu)^{-N}\right)]$, with $\varphi(x_0) = \Delta$, and running N iterations of (GD), the following holds: all iterations belong to the branch $x \geq \tau$, the final iterate is exactly $x_N = \tau$, having the gradient $\nabla\varphi(x_N) = L\tau$ and thus reaching the bound claimed in the first term. □

This one-dimensional worst-case example is inspired from [116, §4.1.2].

Proposition 3.6.3 (Tightness of Theorem 3.2.1 (convex))**.** *Let $L > 0$, $\gamma L \in (0, 2)$ and N be a positive integer. Then there exists x_0 and $f \in \mathcal{F}_{0,L}$ such* *that after applying N iterations of* (GD) *on function f with constant stepsize $\gamma L \in (0,2)$, starting from x_0, it holds:*

$$\tfrac{1}{2L}\|\nabla f(x_N)\|^2 = \frac{f(x_0)-f_*}{\min\left\{1+2N\gamma L\,,\,(1-\gamma L)^{-2N}\right\}}.$$

Proof. With $\Delta > 0$, let $\tau{:=}\sqrt{\frac{2}{L}\,\frac{\Delta}{1+2N\gamma L}}$ and $x_0 = \tau(1{+}N\gamma L)$. We define $\varphi\colon \mathbb{R} \to \mathbb{R}$, such that $\varphi(x_0) = \Delta$, $\varphi(x_*) = 0$ and

$$\varphi(x) := \begin{cases} L\tau|x| - L\frac{\tau^2}{2}, & \text{if } |x| \geq \tau; \\ L\frac{x^2}{2}, & \text{if } |x| < \tau, \end{cases} \qquad \nabla\varphi(x) := \begin{cases} L\tau\,\mathrm{sgn}(x), & \text{if } |x| \geq \tau; \\ Lx, & \text{if } |x| < \tau. \end{cases}$$

Then we select $f = \varphi$ or $f = f_L$ (see Proposition 3.6.1) whether $\gamma L < \overline{\gamma L}_N(0)$ or $\gamma L \geq \overline{\gamma L}_N(0)$, respectively. □

3.6.2 Removing the optimal point in tightness analysis for smooth weakly convex functions

Lemma 3.6.1 shows that by incorporating inequalities (3.20) for all iterates $i \in \mathcal{I}$ the upper bounds remain tight and provides an existence guarantee of an interpolating function with global minimum f_*. It generalizes Theorem 7 from [43], obtained for the particular choices $\mu = -L$ and $\mu = 0$, to any $\mu \leq 0$.

Lemma 3.6.1 (Characterization of the optimal point)**.** *Let $\mathcal{T} = \{(x_i, g_i, f_i)\}_{i\in\mathcal{I}}$ be $\mathcal{F}_{\mu,L}$-interpolable, with $\mu \in (-\infty, 0]$. There exists at least one interpolating function with a finite global minimum*

$$f_* = \min_{x\in\mathbb{R}^d} f(x) = \min_{i\in\mathcal{I}} \left\{f_i - \tfrac{1}{2L}\|g_i\|^2\right\}. \tag{3.34}$$

A global minimizer is given by $x_* = x_{i_*} - \frac{1}{L}g_{i_*}$, where $i_* \in \arg\min_{i\in\mathcal{I}} \left\{f_i - \frac{1}{2L}\|g_i\|^2\right\}$.

Proof. The proof consists on algebraic manipulations, see Section 3.B. □

Lemma 3.6.2 (Decoupling optimal point in tightness analysis)**.** *Assume an $\mathcal{F}_{\mu,L}$-interpolable set $\mathcal{T} = \{(x_i, g_i, f_i)\}_{i\in\mathcal{I}}$ satisfying the condition*

$$N \in \arg\min_{i\in\mathcal{I}} \left\{f_i - \tfrac{1}{2L}\|g_i\|^2\right\}. \tag{3.35}$$

Then proving tightness of the bound

$$\frac{1}{2L}\min_{0\le i\le N}\{\|\nabla f(x_i)\|^2\}\le\frac{f(x_0)-f_*}{1+P_N(\gamma L,\gamma\mu)}$$

reduces to showing tightness of the bound

$$\frac{1}{2L}\min_{0\le i\le N}\{\|\nabla f(x_i)\|^2\}\le\frac{f(x_0)-f(x_N)}{P_N(\gamma L,\gamma\mu)}. \tag{3.36}$$

Proof. Condition (3.35) implies that (3.34) from Lemma 3.6.1 holds with $N=i_*$, hence there exists a function for which $f(x_N)-f_*\ge\frac{1}{2L}\|\nabla f(x_N)\|^2$. This inequality produces the shift in the denominator when appended to rate (3.36). □

Further on, in all tightness proofs we show that condition (3.35) holds and therefore we reduce the analysis to only prove the equality in (3.36).

3.6.3 Tightness analysis for weakly convex functions and stepsizes below $\overline{\gamma L}_1$

In this section, we construct worst-case examples for $\gamma L\le\overline{\gamma L}_1(\kappa)$. For a given choice of parameters, we illustrate these worst-case functions in Figure 3.7.

Range $\gamma_i L\in(0,1]$. We show a one-dimensional worst-case function example which extends to weakly convex functions the one from [2, Proposition 4].

Proposition 3.6.4 (Tightness for $\gamma L\in(0,1]$ (weakly convex))**.** *Let $L>0$, $\mu\in(-\infty,0)$, N be a positive integer and $\gamma_i L\in(0,1]$, with $i=0,\dots,N-1$. Then there exists x_0 and $f\in\mathcal{F}_{\mu,L}$ such that after applying N iterations of (GD) on function f with stepsizes $\gamma_i L$, starting from x_0, it holds:*

$$\tfrac{1}{2L}\min_{0\le i\le N}\{\|\nabla f(x_i)\|^2\}=\frac{f(x_0)-f_*}{1+\sum_{i=0}^{N-1}\gamma_i L\big(2+\frac{\gamma_i L\gamma_i\mu}{\gamma_i L-\gamma_i\mu}\big)}.$$

Proof. Motivated by the necessary conditions $g_i=g_{i+1}$ in Lemma 3.5.1, a one-dimensional function with constant gradients U at the iterates is built, where U is the square root of the minimum gradient norm: $U^2:=\frac{\Delta}{\frac{1}{2L}\sum_{i=0}^{N-1}\gamma_i L(2+\frac{\gamma_i L\gamma_i\mu}{\gamma_i L-\gamma_i\mu})}$, with $\Delta > 0$. The equality case in Lemma 3.5.1 implies (together with $g_i = U, \forall i = 0, 1, \ldots, N$):

$$\frac{f_i - f_{i+1}}{\Delta} = \frac{\gamma_i L\big(2 + \frac{\gamma_i L \gamma_i \mu}{\gamma_i L - \gamma_i \mu}\big)}{\sum_{j=0}^{N-1} \gamma_j L\big(2 + \frac{\gamma_j L \gamma_j \mu}{\gamma_j L - \gamma_j \mu}\big)} \iff \frac{f_i}{\Delta} = \frac{\sum_{j=i}^{N-1} \gamma_j L\big(2 + \frac{\gamma_j L \gamma_j \mu}{\gamma_j L - \gamma_j \mu}\big)}{\sum_{j=0}^{N-1} \gamma_j L\big(2 + \frac{\gamma_j L \gamma_j \mu}{\gamma_j L - \gamma_j \mu}\big)}.$$

Without loss of generality, we set $f_N = 0$, $f_0 = \Delta$ and $x_N = 0$ and define the triplets $\big\{(x_i, g_i, f_i)\big\}_{i \in \{0,\ldots,N\}}$, with $g_i = U$ and $x_i = U \sum_{j=i}^{N-1} \gamma_j$. Consider the following points lying between consecutive iterates:

$$\bar{x}_i \coloneqq x_i - \tfrac{-\mu}{L-\mu} \gamma_i U \in [x_{i+1}, x_i], \quad i = 0, 1, \ldots, N-1.$$

A worst-case function interpolating the triplets is the piecewise quadratic $f \colon \mathbb{R} \to \mathbb{R}$, with alternating curvature between μ and L with inflection points x_i and $\bar{x}_i$, constant gradients at the iterations and decreasing function values, where $i \in \{0, \ldots, N-1\}$:

$$f(x) = \begin{cases} \frac{L}{2}(x - x_N)^2 + U(x - x_N) + f_N, & \text{if } x \in (-\infty, x_N]\,; \\ \frac{\mu}{2}\left(x - x_{i+1}\right)^2 + U\left(x - x_{i+1}\right) + f_{i+1}, & \text{if } x \in [x_{i+1}, \bar{x}_i]; \\ \frac{L}{2}\left(x - x_i\right)^2 + U\left(x - x_i\right) + f_i, & \text{if } x \in [\bar{x}_i, x_i]; \\ \frac{L}{2}\left(x - x_0\right)^2 + U\left(x - x_0\right) + f_0, & \text{if } x \in [x_0, \infty)\,. \end{cases} \tag{3.37}$$

□

An illustration of the one-dimensional worst-case construction from Proposition 3.6.4 is given in Figure 3.7a.

Corollary 3.6.1 (Tightness of Corollary 3.2.2 (convex))**.** *Let $L > 0$, N be a positive integer and $\gamma_i L \in (0, \frac{3}{2}]$, with $i = 0, \ldots, N-1$. Then there exists x_0 and $f \in \mathcal{F}_{0,L}$ such that after applying N iterations of* (GD) *on function f with stepsizes $\gamma_i L$, starting from x_0, it holds:*

$$\tfrac{1}{2L}\|\nabla f(x_N)\|^2 = \frac{f(x_0) - f_*}{1 + 2\sum_{i=0}^{N-1} \gamma_i L}.$$

Proof. The proof results by setting $\mu = 0$ in Proposition 3.6.4 to obtain a linear function between iterations, extended outside by a quadratic of curvature L. □

Range $\gamma_i L \in (1, \overline{\gamma L}_1(\kappa)]$. We build a set of triplets matching the corresponding bounds from Theorems 3.2.3 and 3.2.4 and interpolating an $\mathcal{F}_{\mu,L}$-function, implying the validity of all interpolation inequalities relating the $N+1$ iterations. The upper bounds in Theorem 3.2.4 are obtained by analyzing only consecutive iterations, from which we track equality conditions collected in Proposition 3.6.5.

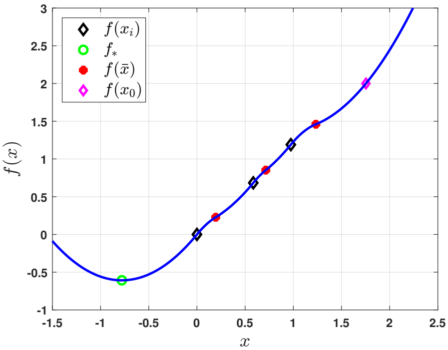

(a) One-dimensional worst-case function example corresponding to short stepsizes $\gamma_i L \in (0,1]$ (piecewise quadratic; see Proposition 3.6.4). Setup: $N=3$, $f_0 - f_* = 2$, $L=2$, $\mu=-4$, $\gamma_0 L = 1$, $\gamma_1 L = 0.5$, $\gamma_2 L = 0.75$.

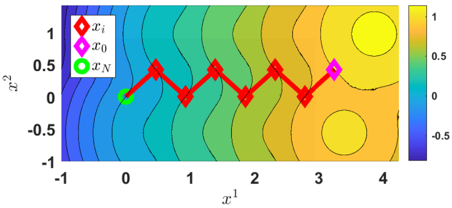

(b) Function example for constant stepsize $\gamma L \in (1, \overline{\gamma L}_1(\kappa)]$. Setup: $L=1$, $\mu=-0.5$, $\gamma L = 1.5$, $N=7$. Component x^1 is successively reduced, while x^2 alternates between two values, due to alternating gradients g_i and g_{i+1}.

Figure 3.7: Worst-case examples for weakly convex functions and stepsizes below the first stepsize threshold ($\gamma L \leq \overline{\gamma L}_1(\kappa)$), matching the bounds from Theorems 3.2.3 and 3.2.4.

Proposition 3.6.5 (Necessary conditions for worst case with variable stepsizes (weakly convex))**.** *Let $L>0$, $\mu \in (-\infty, 0)$, N be a positive integer and $\gamma_i L \in (1, \overline{\gamma L}_1(\kappa))$, with $i = 0, \ldots, N-1$. Then any worst-case function $f \in \mathcal{F}_{\mu,L}$ matching the bound from Theorem 3.2.4 satisfies the following identities on the triplets $\mathcal{T} = \{(x_i, g_i, f_i)\}_{i \in \{0,\ldots,N\}}$, with $x_{i+1} = x_i - \gamma_i g_i$:*

$$\|g_i\|^2 = U^2 := \frac{f_0 - f_N}{\frac{1}{2L}\sum_{j=0}^{N-1} \frac{\gamma_j L(2-\gamma_j L)(2-\gamma_j \mu)}{2-\gamma_j L - \gamma_j \mu}}, \quad \forall i = 0, \ldots, N; \tag{3.38}$$

$$c(\gamma_i L, \gamma_i \mu) := \frac{\langle g_i, g_{i+1} \rangle}{U^2} = \frac{1 + (1-\gamma_i \mu)(1-\gamma_i L)}{2 - \gamma_i L - \gamma_i \mu}, \quad \forall i = 0, \ldots, N-1; \tag{3.39}$$

$$f_i = f_N + \frac{U^2}{2L} \sum_{j=i}^{N-1} \frac{\gamma_j L(2-\gamma_j \mu)(2-\gamma_j L)}{2-\gamma_j L - \gamma_j \mu}, \quad \forall i = 0, \ldots, N. \tag{3.40}$$

In particular, a two-dimensional function $f\colon \mathbb{R}^2 \to \mathbb{R}$ satisfies these necessary conditions if, for any suitable g_0, the following holds:

$$g_{i+1} = R\big((-1)^{i+1}\theta(\gamma_i L, \gamma_i \mu)\big) g_i, \quad \forall i = 0, 1, \ldots, N-1, \tag{3.41}$$

where R is the two-dimensional rotation matrix and

$$\theta(\gamma_i L, \gamma_i \mu) \coloneqq \arccos\big(c(\gamma_i L, \gamma_i \mu)\big). \tag{3.42}$$

Moreover,

$$N \in \underset{0\le i\le N}{\arg\min}\left\{f_i - \tfrac{1}{2L}\|g_i\|^2\right\}. \tag{3.43}$$

Proof. See Section 3.C.1. □

Proposition 3.6.6 (Tightness for $\gamma L \in (1, \overline{\gamma L}_1]$ (weakly convex))**.** *Let $L > 0$, $\mu \in (-\infty, 0)$, N be a positive integer and $\gamma L \in (1, \overline{\gamma L}_1(\kappa)]$. Then there exists x_0 and $f \in \mathcal{F}_{\mu,L}$ such that after applying N iterations of (GD) on function f with constant stepsize γL, starting from x_0, it holds:*

$$\tfrac{1}{2L} \min_{0\le i\le N}\{\|\nabla f(x_i)\|^2\} = \frac{f(x_0) - f_*}{1 + \gamma L \frac{(2-\gamma L)(2-\gamma\mu)}{2-\gamma L-\gamma\mu} N}.$$

Proof. We construct a two-dimensional worst-case function $f : \mathbb{R}^2 \to \mathbb{R}$. From Proposition 3.6.5, the rotation angle between the gradients $\theta(\gamma L, \gamma\mu)$ is constant, hence they alternate between odd and even iterations such that:

$$\|g_i\|^2 = \langle g_i, g_{i+(2\mathbb{N})}\rangle = U^2 \ \text{ and } \ \langle g_i, g_{i+(2\mathbb{N}+1)}\rangle = cU^2 \qquad \forall i = 0, \dots, N-1.$$

Consider the set of triplets $\mathcal{T}_\mathcal{I} \coloneqq \{(x_i, g_i, f_i)\}_{i\in\{0,\dots,N\}}$ defined as follows: (i) since g_0 can be set (almost) arbitrary, we take $g_i = U\left[\sqrt{\frac{1+c}{2}};\, (-1)^i\sqrt{\frac{1-c}{2}}\right]^\top$; (ii) by setting $x_N = [0, 0]^\top$, the iterations are given by $x_i = \gamma\sum_{j=i}^{N-1} g_j$; (iii) by setting $f_N = 0$, the function values simplify to $f_i = f_0(1 - \frac{i}{N})$. In Section 3.C.1 we demonstrate that $\mathcal{T}_\mathcal{I}$ is $\mathcal{F}_{\mu,L}$-interpolable. □

Figure 3.7b shows a numerical example of the triplets from Proposition 3.6.6's proof. Excepting initial point x_0, all subsequent iterations are inflection points where the curvature transitions from μ to L. This pattern is reversed between odd and even iterations and the decrease mainly occurs to component x^1.

Remark 3.6.2 (QCQP-based numerical worst-case construction)**.** *The numerical example from Figure 3.7b is obtained from the lower interpolation function derived in [110, Theorem 3.14], originally defined for smooth strongly convex functions. This interpolation function can be extended to arbitrary lower curvatures μ using the minimal curvature subtraction argument employed in the proof of Theorem 2.2.1 (refer to Remark 3.13 in [110]). Given* *triplets $\mathcal{T}_{\mathcal{I}} := \{(x_i, g_i, f_i)\}_{i\in\mathcal{I}}$, this function is the solution of the quadratically constrained optimization problem (QCQP):*

$$f(x) := \min_{g\in\mathbb{R}^2} \max_{0\leq i\leq N} \left\{ f_i + \langle g_i\,,\, x - x_i\rangle + \tfrac{\mu}{2}\|x-x_i\|^2 + \tfrac{1}{2(L-\mu)}\|g - g_i - \mu(x-x_i)\|^2 \right\}.$$

When using non-constant stepsizes $\gamma_i L \in (1, \overline{\gamma L}_1(\kappa)]$, based on the triplets construction from Proposition 3.6.5 we state in Conjecture 3.6.1 the tightness of the convergence rate from Theorem 3.2.4. The conjecture relies on numerical verification that all interpolation conditions hold (see Footnote 2).

Conjecture 3.6.1 (Tightness for $\gamma_i L \in (1, \overline{\gamma L}_1]$ (weakly convex))**.** *There exists a two-dimensional function $f \in \mathcal{F}_{\mu,L}$ interpolating the triplets defined in Proposition 3.6.5, satisfying*

$$\tfrac{1}{2L} \min_{0\leq i\leq N}\{\|\nabla f(x_i)\|^2\} = \frac{f(x_0) - f(x_N)}{\sum_{i=0}^{N-1} \frac{\gamma_i L(2-\gamma_i L)(2-\gamma_i\mu)}{2-\gamma_i L-\gamma_i\mu}}.$$

3.6.4 Tightness analysis for weakly convex functions and stepsizes above $\overline{\gamma L}_1$

We firstly show the tightness for the stepsize interval $[\overline{\gamma L}_{N-1}(\kappa), \overline{\gamma L}_N(\kappa))$ in Proposition 3.6.7, by constructing a two-dimensional quadratic function. Then Proposition 3.6.8 demonstrates the exactness of our bounds for any stepsize in the range $(1, \overline{\gamma L}_N(\kappa))$, by providing in its proof a three-dimensional example. We also provide alternative, equivalent expressions of denominator $P_N(\gamma L, \gamma\mu)$ from (3.6). Using $\bar{N}(\gamma L, \kappa)$ (recall Definition 3.5.1) emphasizes that the analysis of worst-case functions relies on fixing the stepsizes and varying the number of iterations; $\bar{N}(\gamma L, \kappa)$ plays the role of index k delimiting intervals in the proofs of Theorem 3.2.3 (see Section 3.5.4) and Theorem 3.2.2 (see Section 3.5.5).

Proposition 3.6.7 (Tightness for $\gamma L \in [\overline{\gamma L}_{N-1}, \overline{\gamma L}_N)$ (weakly convex))**.** *Let $L > 0$, $\mu \in (-\infty, 0)$, N be a positive integer and $\gamma L \in [\overline{\gamma L}_{N-1}(\kappa), \overline{\gamma L}_N(\kappa))$, such that $\bar{N}(\gamma L, \kappa) = N - 1$. Then there exists x_0 and $f \in \mathcal{F}_{\mu,L}$ such that after applying N iterations of* (GD) *on function f with constant stepsize γL, starting from x_0, the corresponding performance bound from Theorem 3.2.3 holds exactly:*

$$\tfrac{1}{2L} \min_{0\leq i\leq N}\{\|\nabla f(x_i)\|^2\} = \frac{f(x_0) - f_*}{1 + \gamma L \frac{\frac{(1-\gamma L)^2\, E_N(1-\gamma L)}{1-(1-\gamma L)^2} - \frac{(1-\gamma\mu)^2\, E_N(1-\gamma\mu)}{1-(1-\gamma\mu)^2}}{\frac{1}{1-(1-\gamma L)^2} - \frac{1}{1-(1-\gamma\mu)^2}}},$$

with $E_k(x)$ given in Definition 3.2.1.

Proof. Let $\Delta > 0$ and $U = \sqrt{\frac{2L\,\Delta}{\gamma L P_N(\gamma L,\gamma\mu)}}$, with $P_N(\gamma L,\gamma\mu)$ defined in (3.6). The following performance bound (without f_*)

$$\frac{1}{2L}\min_{0\le i\le N}\{\|\nabla f(x_i)\|^2\} = \frac{f(x_0)-f(x_N)}{\gamma L}\frac{\frac{1}{1-(1-\gamma L)^2}-\frac{1}{1-(1-\gamma\mu)^2}}{\frac{(1-\gamma L)^2\,E_N(1-\gamma L)}{1-(1-\gamma L)^2}-\frac{(1-\gamma\mu)^2\,E_N(1-\gamma\mu)}{1-(1-\gamma\mu)^2}}$$

is achieved by the $\mathcal{F}_{\mu,L}$ quadratic function

$$f(x) = \frac{1}{2}(x-x_N)^\top \begin{bmatrix} L & 0\\ 0 & \mu\end{bmatrix}(x-x_N) + g_N^\top(x-x_N) + f_N,$$

with arbitrarily fixed $x_N = [0,0]^\top$, $f_N = 0$, $f(x_0) = \Delta$ and

$$g_N \coloneqq U\sqrt{\frac{[(1-\gamma\mu)^2-1][1-(1-\gamma L)^2]}{(1-\gamma\mu)^2-(1-\gamma L)^2}}\begin{bmatrix}\frac{1-\gamma L}{\sqrt{1-(1-\gamma L)^2}}\\ \frac{1-\gamma\mu}{\sqrt{(1-\gamma\mu)^2-1}}\end{bmatrix}.$$

□

Remark 3.6.3 (Componentwise gradient scaling factors)**.** *The gradients of the worst-case function from Proposition 3.6.7, evaluated at iterations x_i, decrease in each component by the scaling factors $(1-\gamma L)$ and $(1-\gamma\mu)$, respectively:*

$$\nabla f(x_{i+1}) = \begin{bmatrix} 1-\gamma L & 0\\ 0 & 1-\gamma\mu\end{bmatrix}\nabla f(x_i).$$

Proving the tightness of Theorem 3.2.3 for stepsizes $\gamma L \in \left[\overline{\gamma L}_1(\kappa), \overline{\gamma L}_{N-1}(\kappa)\right)$ presents a significant challenge due to the nonlinear term in (3.6), which connects the quadratic function characterizing the first $\bar{N}(\gamma L,\kappa)+1$ iterations (see Proposition 3.6.7) and the 2D function with alternating gradients directions from Proposition 3.6.6. This later function dominates the sublinear rate through the leading coefficient $\gamma L\, p(\gamma L,\gamma\mu)$. The three-dimensional worst-case construction from Proposition 3.6.8's proof also provides an alternative worst-case construction for this regime with stepsizes $\gamma L \in (1, \overline{\gamma L}_1]$.

Proposition 3.6.8 (Tightness for $\gamma L \in (1,\overline{\gamma L}_N)$ (weakly convex))**.** *Let $L>0$, $\mu \in (-\infty, 0)$, N be a positive integer and $\gamma L \in (1, \overline{\gamma L}_N(\kappa))$. There exists x_0 and $f \in \mathcal{F}_{\mu,L}$ such that after applying N iterations of (GD) on function f with constant stepsize γL, starting from x_0, the performance bound from Theorem 3.2.3 is attained:*

$$\frac{1}{2L}\min_{0\le i\le N}\{\|\nabla f(x_i)\|^2\} = \frac{f(x_0)-f_*}{1+\frac{(2-\gamma L)(2-\gamma\mu)}{2-\gamma L-\gamma\mu}(N-\bar{N}) + \frac{\frac{E_{\bar{N}}(1-\gamma L)}{1-(1-\gamma L)^2}-\frac{E_{\bar{N}}(1-\gamma\mu)}{1-(1-\gamma\mu)^2}}{\frac{1}{1-(1-\gamma L)^2}-\frac{1}{1-(1-\gamma\mu)^2}}},$$

where $E_k(x)$ given in Definition 3.2.1, together with the complementary rate:

$$\frac{1}{2L}\min_{0\le i\le N}\{\|\nabla f(x_i)\|^2\} = \frac{f(x_0)-f(x_N)}{\frac{(2-\gamma L)(2-\gamma\mu)}{2-\gamma L-\gamma\mu}(N-\bar{N}) + \frac{\frac{E_{\bar{N}}(1-\gamma L)}{1-(1-\gamma L)^2}-\frac{E_{\bar{N}}(1-\gamma\mu)}{1-(1-\gamma\mu)^2}}{\frac{1}{1-(1-\gamma L)^2}-\frac{1}{1-(1-\gamma\mu)^2}}}.$$

Proof. Let $\Delta > 0$ and $U = \sqrt{\frac{2L\,\Delta}{\gamma L P_N(\gamma L,\gamma\mu)}}$, with $P_N(\gamma L,\gamma\mu)$ defined in (3.6). The value of U denotes the minimum gradient norm over N iterations, i.e., the rate, where Δ accounts for the initial function value gap $f(x_0)-f(x_N)$. The definition of $\bar{N}=\bar{N}(\gamma L,\kappa)$ implies $\gamma L \in [\overline{\gamma L}_{\bar{N}-1}(\kappa), \overline{\gamma L}_{\bar{N}}(\kappa))$, with $1 \le \bar{N} \le N-1$. We define the set of triplets $\mathcal{T}_{\mathcal{I}} := \{(x_i, g_i, f_i)\}_{i=\{0,1,\dots,N\}}$ as follows: without loss of generality, we fix $f_N = 0$ and $x_N = 0$ and consider

$$g_i = U\sqrt{\frac{[(1-\gamma\mu)^2-1][1-(1-\gamma L)^2]}{(1-\gamma\mu)^2-(1-\gamma L)^2}}\begin{bmatrix}\frac{(1-\gamma L)^{i-\bar{N}}}{\sqrt{1-(1-\gamma L)^2}}\\ \frac{(1-\gamma\mu)^{i-\bar{N}}}{\sqrt{(1-\gamma\mu)^2-1}}\\ 0\end{bmatrix} \quad \forall i = 0,1,\dots,\bar{N}+1;$$

$$g_{i+1} = (1-\gamma L)g_i + \gamma(L-\mu)\cos(\theta_{i+1})\begin{bmatrix}0 & \\ & R(\theta_{i+1})\end{bmatrix}g_i \quad \forall i = \bar{N},\dots,N;$$

$$q_i = (-1)^{i-(\bar{N}+1)}\sqrt{[(1-\gamma\mu)^2-1]\frac{1-(1-\gamma L)^{2(i-(\bar{N}+1))}}{1-(1-\gamma L)^2}} \quad \forall i = \bar{N}+1,\dots,N;$$

$$\theta_i = \arctan(q_i) \quad \forall i = \bar{N}+1,\dots,N;$$

$$f_i = \Delta\frac{N-i+\frac{\gamma L\gamma\mu}{\gamma L-\gamma\mu}\sum_{j=1}^{\bar{N}-i}T_j(\gamma L,\kappa)}{N+\frac{\gamma L\gamma\mu}{\gamma L-\gamma\mu}\sum_{j=1}^{\bar{N}}T_j(\gamma L,\kappa)} \quad i = 0,1,\dots,N,$$

and $x_{i+1} = x_i - \gamma g_i$, $i = 0,\dots,N-1$ and $R(\theta) := \begin{bmatrix}\cos(\theta) & -\sin(\theta)\\ \sin(\theta) & \cos(\theta)\end{bmatrix}$ being the two-dimensional rotation matrix of angle θ. For $\mathcal{T}_{\mathcal{I}}$, we show in Section 3.C.2 that: (i) it attains the performance bound from Theorem 3.2.3 and (ii) it is interpolable by a $\mathcal{F}_{\mu,L}$-function with three variables. □

Remark 3.6.4 (Construction of the 3D worst-case example)**.** *If the stepsize is not the threshold $\overline{\gamma L}_{\bar{N}}(\kappa)$, then the first $\bar{N}+1$ gradients have uniquely defined norm and inner products between them. This follows the equality conditions from Lemma 3.5.6. Restricting them to be two-dimensional, they are unique up to applying orthonormal transformations. The first $\bar{N}+1$ iterations are exactly characterized by the worst-case function from Proposition 3.6.7, where a third component is set to zero. For the subsequent iterations, the first component* *remains linearly decreasing in $(1-\gamma L)$, while the others are obtained by applying a certain rotation which, given the decrease in the first component, preserves: (i) the gradient norm to U^2 and (ii) the inner product between consecutive iterations to cU^2 such that:*

$$\text{(i): } \|g_{i+1} - (1-\gamma L)g_i\|^2 = U^2\big[1 + (1-\gamma L)^2 - 2(1-\gamma L)c\big];$$

$$\text{(ii): } \langle g_{i+1} - (1-\gamma L)g_i \,,\, g_i\rangle = U^2\big[c - (1-\gamma L)\big],$$

with $c = \frac{1+(1-\gamma L)(1-\gamma\mu)}{2-\gamma L-\gamma\mu}$ from Proposition 3.6.6. The matrix $R(\theta_{i+1})$ rotates the second and third components of g_i by angle θ_i, whose absolute value monotonically increasing towards $\lim_{i\to\infty}|\theta_i| = \arctan(\sqrt{\frac{-\mu}{L}\frac{2-\gamma\mu}{2-\gamma L}})$. These components are then properly scaled to preserve $\|g_{i+1}\| = U$. Moreover, as i approaches ∞, the first component vanishes and the gradients start alternating, resembling the same pattern observed for the two-dimensional worst-case described in Proposition 3.6.6's proof.

3.7 Conclusion

We have presented a comprehensive worst-case analysis of gradient descent applied to smooth weakly convex, convex and strongly convex functions across the entire range of constant stepsizes $\gamma L \in (0, 2)$.

For convex and strongly convex functions we establish and prove tight performance bounds over ranges previously only covered by conjectures.

In the weakly convex scenario, determining convergence rates is more challenging outside extreme cases already addressed in the literature: $\mu \in \{-\infty, 0\}$. The state-of-the-art analysis of smooth not-necessarily convex functions is extended from the particular case $\mu = -L$ and stepsizes $\gamma L \in (0, \sqrt{3}]$ to encompass full domains.

The proofs are based on incorporating not only consecutive interpolation inequalities but also those connecting iterations at a distance of two, the key idea being detailed in Remark 3.5.2. A key advantage of these proofs type is the obtaining of new descent lemma type inequalities, independent of the number of iterations N, as they would be in tight convergence analyses using inequalities involving the last iterate. When deriving closed-form expressions for tight performance analysis with constant stepsizes, it is uncommon and challenging to involve nonconsecutive iterations.

While all rates align with the numerical solutions obtained from solving performance estimation problems (PEPs), we demonstrate analytically their tightness over the complete stepsize range $\gamma L \in (0, 2)$.

This analysis yields insights into the optimal constant stepsizes with respect to the demonstrated worst-case scenarios and we further improve upon them by proposing dynamic stepsize schedules which have the key benefit of being independent of the number of iterations. Testing them on practical problems is left as future work.

Appendix

3.A Properties of stepsize sequences (proofs)

3.A.1 Stepsize thresholds

Notation. In the proofs within this section we use the notation $T_k(\gamma L, \kappa) = \sum_{i=1}^{k} \delta_i(\gamma L, \kappa)$, where

$$\delta_i(\gamma L, \kappa) := \frac{2 - \kappa\gamma L}{(1 - \kappa\gamma L)^{2i}} - \frac{2 - \gamma L}{(1 - \gamma L)^{2i}}.$$

Proof of Proposition 3.2.1. We show $\frac{dT_k(\gamma L,\kappa)}{d(\gamma L)} \geq 0$ for $\gamma L \in (1, \overline{\gamma L}_\infty(\kappa))$, which resumes to proving the positivity of

$$\frac{d\delta_i(\gamma L, \kappa)}{d(\gamma L)} := \frac{-\kappa\left(1 - 2i\frac{2-\kappa\gamma L}{1-\kappa\gamma L}\right)}{(1 - \kappa\gamma L)^{2i}} + \frac{1 + 2i\frac{2-\gamma L}{\gamma L - 1}}{(1 - \gamma L)^{2i}}.$$

1. **Case $\kappa \geq 0$:** we lower bound $\frac{d\delta_i(\gamma L,\kappa)}{d(\gamma L)}$ by neglecting the denominators:

$$\delta_i(\gamma L, \kappa) \geq -\kappa\Big(1 - 2i\frac{2 - \kappa\gamma L}{1 - \kappa\gamma L}\Big) + 1 + 2i\frac{2 - \gamma L}{\gamma L - 1}$$

$$= (1 - \kappa) + 2i\Big[\frac{\kappa(2 - \kappa\gamma L)}{1 - \kappa\gamma L} + \frac{2 - \gamma L}{\gamma L - 1}\Big] \geq 0.$$

2. **Case $\kappa < 0$:** it is sufficient to show

$$\frac{1 + 2i\frac{2-\gamma L}{\gamma L-1}}{(1 - \gamma L)^{2i}} \geq 1 \geq -\kappa\frac{-1 + 2i(1 + \frac{1}{1-\kappa\gamma L})}{(1 - \kappa\gamma L)^{2i}}.$$

The l.h.s. holds because $\gamma L \in (1,2)$ and $(1-\gamma L)^{2i} < 1$. The r.h.s. is equivalent to:

$$(1-\kappa\gamma L)^{2i} \geq -\kappa\Big[-1+2i\Big(1+\frac{1}{1-\kappa\gamma L}\Big)\Big].$$

Using $\gamma L \in (1,2)$ and the second order Taylor approximation for $(1-\kappa)^{2i}$, we get the bounds:

$$(1-\kappa\gamma L)^{2i} \geq (1-\kappa)^{2i} \geq 1+2i(-\kappa)+i(2i-1)\kappa^2$$

$$-\kappa\Big[-1+2i\Big(1+\frac{1}{1-\kappa\gamma L}\Big)\Big] \leq -\kappa\Big[-1+2i\left(1+\frac{1}{1-\kappa}\right)\Big].$$

Therefore, it is sufficient to show:

$$1+2i(-\kappa)+i(2i-1)\kappa^2 \geq -\kappa\Big[-1+2i\Big(1+\frac{1}{1-\kappa}\Big)\Big],$$

which equivalently rewrites as the following inequality valid for all $i \geq 1$:

$$\kappa^2 i(i-1)+\Big(\kappa i+\frac{1}{1-\kappa}\Big)^2+\frac{(1-\kappa)^3-1}{(1-\kappa)^2} \geq 0.$$

□

Lemma 3.A.1 establishes the existence of an index beyond which the contribution δ becomes negative. This observation is key to proving Proposition 3.2.2.

Lemma 3.A.1. *Let $\kappa \in (-\infty, 1)$ and $\gamma L \in (1, \overline{\gamma L}_\infty(\kappa))$. Then there exists an index k such that $\delta_k(\gamma L, \kappa) < 0$ and $\delta_j(\gamma L, \kappa) < 0$ for all $j > k$.*

Proof. The sequence δ_i is a difference of two monotone sequences: $\frac{2-\kappa\gamma L}{(1-\kappa\gamma L)^{2i}}$ and $\frac{2-\kappa\gamma L}{(1-\kappa\gamma L)^{2i}}$, respectively. The former is decreasing for $\kappa < 0$ (weakly convex functions) and increasing for $\kappa \geq 0$ (convex functions), whereas the later is always decreasing. In particular, $\delta_0(\gamma L, \kappa) > 0$. Because $(1-\gamma\mu)^2 \leq (1-\gamma L)^2$, to the limit it holds $\lim_{k\to\infty} \delta_k(\gamma L, \kappa) = -\infty$ and there exists some index k with $\delta_k(\gamma L, \kappa) < 0$. We prove by contradiction that $\delta_j(\gamma L, \kappa) < 0$ for all $j > k$; assume there exists some index $l > k$ such that $\delta_l(\gamma L, \kappa) \geq 0$. Then the following inequalities would hold:

$$\delta_k < 0: \ \frac{2-\gamma\mu}{(1-\gamma\mu)^{2k}} < \frac{2-\gamma L}{(1-\gamma L)^{2k}};$$

$$\delta_l > 0: \ \frac{2-\gamma\mu}{(1-\gamma\mu)^{2l}} > \frac{2-\gamma L}{(1-\gamma L)^{2l}}.$$

By multiplying the inequalities and performing the simplifications we get:

$$(1-\gamma\mu)^{2(k-l)} > (1-\gamma L)^{2(k-l)},$$

which implies, due to $(1-\gamma\mu)^2 > (1-\gamma L)^2$, that $k > l$, contradiction! □

Proof of Proposition 3.2.2. We prove these claims in turn.

1. Case $k = 0$ holds by definition since $\overline{\gamma L}_0(\kappa) = 1$. For any $k \geq 1$, by using Lemma 3.A.1, condition $T_k(\overline{\gamma L}_k, \kappa) = 0$ implies $\delta_k(\gamma L, \kappa) \leq 0$ and hence $\delta_{k+1}(\gamma L, \kappa) < 0$. Furthermore, $T_{k+1}(\overline{\gamma L}_k, \kappa) \leq T_k(\overline{\gamma L}_k, \kappa) = 0$. Then it holds:

$$T_{k+1}(\overline{\gamma L}_k, \kappa) \leq T_k(\overline{\gamma L}_k, \kappa) = 0 = T_{k+1}(\overline{\gamma L}_{k+1}, \kappa).$$

The monotonicity of $T_k(\gamma L, \kappa)$ in stepsize γL (see Proposition 3.2.1) implies $\overline{\gamma L}_k(\kappa) < \overline{\gamma L}_{k+1}(\kappa)$.
2. Because stepsize thresholds sequence is monotone, it has a limit.
(a) **Case $\kappa > 0$.** By definition, $\overline{\gamma L}_k(\kappa) < \frac{2}{1+\kappa} < 2$ for all integers $k \geq 1$ and condition $T_k(\overline{\gamma L}_k(\kappa), \kappa) = 0$ implies:

$$\left(\frac{1-\overline{\gamma L}_k(\kappa)}{1-\kappa\overline{\gamma L}_k(\kappa)}\right)^{2k} = \kappa + (1-\kappa)(1-\overline{\gamma L}_k(\kappa))^{2k} \iff$$

$$\left|\frac{1-\overline{\gamma L}_k(\kappa)}{1-\kappa\overline{\gamma L}_k(\kappa)}\right| = \exp\left[\frac{\ln\left(\kappa + (1-\kappa)(1-\overline{\gamma L}_k(\kappa))^{2k}\right)}{2k}\right].$$

To the limit $k \to \infty$, the r.h.s. equals 1 because $(1-\overline{\gamma L}_k(\kappa))^{2k} \to 0$, hence $|1-\overline{\gamma L}_k(\kappa)| = |1-\kappa\overline{\gamma L}_k(\kappa)|$, which implies the solution $\overline{\gamma L}_\infty(\kappa) = \frac{2}{1+\kappa}$.

(b) **Case $\kappa < 0$.** By definition, $\overline{\gamma L}_k(\kappa) < 2$ for all integers $k \geq 1$ and condition $T_k(\overline{\gamma L}_k(\kappa), \kappa) = 0$ implies:

$$(1-\overline{\gamma L}_k(\kappa))^{-2k} = 1 + \frac{1-(1-\kappa\overline{\gamma L}_k(\kappa))^{-2k}}{-\kappa} \iff$$

$$\left|1-\overline{\gamma L}_k(\kappa)\right| = \exp\left[\frac{-\ln\left(1+\frac{1-(1-\kappa\overline{\gamma L}_k(\kappa))^{-2k}}{-\kappa}\right)}{2k}\right].$$

To the limit $k \to \infty$, the r.h.s. equals 1 because $(1-\kappa\overline{\gamma L}_k(\kappa))^{-2k} \to 0$ due to $\kappa < 0$. Hence, $\left|1-\overline{\gamma L}_k(\kappa)\right| = 1$, with the solution $\overline{\gamma L}_\infty(\kappa) = 2$.
(c) **Case $\kappa = 0$.** By definition, $\overline{\gamma L}_k(0) < 2$ and, similarly to the case $\kappa < 0$:

$$\left|1-\overline{\gamma L}_k(\kappa)\right| = \left[1+2k\overline{\gamma L}_k(\kappa)\right]^{\frac{-1}{2k}} = \exp\left[\frac{-\ln\left(1+2k\overline{\gamma L}_k(\kappa)\right)}{2k}\right].$$

To the limit $k \to \infty$, the r.h.s. equals 1, hence $\left|1 - \overline{\gamma L}_k(\kappa)\right| = 1$, with the solution $\overline{\gamma L}_\infty(\kappa) = 2$.
3. Condition $\gamma L \in [\overline{\gamma L}_k(\kappa), \overline{\gamma L}_{k+1}(\kappa))$ implies $T_k(\overline{\gamma L}_k(\kappa), \kappa) \geq 0$ and $T_{k+1}(\overline{\gamma L}_k(\kappa), \kappa) < 0$. Because $\delta_i(\gamma L, \kappa) < 0$ for all $i \geq k$, we have (i) $T_i(\overline{\gamma L}_k, \kappa) > 0$ for $i = 1, \ldots, k-1$ (since k is uniquely defined) and (ii) $T_i(\overline{\gamma L}_k, \kappa) < 0$ for $i \geq k+1$.

□

3.A.2 Proof of optimal constant stepsize for weakly convex functions

Proof of Proposition 3.3.3. Minimizing the upper bound from Theorem 3.2.3 is equivalent to maximizing the denominator $P_N(\gamma L, \gamma\mu)$ from (3.6). For a fixed curvature ratio $\kappa = \frac{\mu}{L}$, this reduces to maximizing the quantity $r(l) := l\, p(l, \kappa l)$, where $l \in (0, 2)$ is the normalized stepsize γL and $p(l, u)$ is defined in (3.7):

$$r(l) = l\left[2 - \frac{-\kappa l^2}{1-\kappa l - |1-l|}\right] = \begin{cases} l\left(2 - \frac{\kappa l^2}{l-\kappa l}\right) & l \in (0, 1]; \\ l\,\frac{(2-l)(2-\kappa l)}{2-l-\kappa l} & l \in [1, \overline{\gamma L}_1(\kappa)]. \end{cases}$$

Its first branch is concave in l and reaches its maximum for $l = 1$. Since this value is also feasible for the second branch, it means that the maximum belongs to the later expression and restrict the analysis to the interval $[1, 2)$. One can check that $r(l)$ is strictly concave for $l \in [1, 2)$. Moreover, $r'(1) = \frac{2}{(1-\kappa)^2} > 0$ and $r'(2) = -2(1 - \frac{1}{\kappa}) < 0$, where

$$r'(l) = \frac{2}{(2-l-\kappa l)^2}\left[-\kappa\big(1+\kappa\big)l^3 + \left[3\kappa + (1+\kappa)^2\right]l^2 - 4\big(1+\kappa\big)l + 4\right]. \tag{3.44}$$

Hence, there exists a unique maximizer $l^* \in [1, 2)$, further on denoted by $(\gamma L)^*(\kappa)$. The square bracket from (3.44) is exactly the expression from equation (3.14).

Moreover, the optimal stepsize $(\gamma L)^*(\kappa)$ strictly increases with κ, from $(\gamma L)^*(-\infty) = 1$ to $(\gamma L)^*(0) \nearrow 2$ and crosses the threshold $\overline{\gamma L}_1(\kappa)$ at some specific $\bar{\kappa}$ (see Figure 3.1), which results by taking $(\gamma L)^*(\kappa) = \overline{\gamma L}_1(\kappa)$ in (3.14). For longer stepsizes, some other transient and exponential terms in γ are involved (see expression of denominator (3.6)), therefore the solution $(\gamma L)^*(\kappa) > \overline{\gamma L}_1(\kappa)$ is only asymptotically optimal ($N \to \infty$). □

3.A.3 Scheduled stepsizes

Proposition 3.3.4. We present the proofs for each statement:

(i) One can check that $s_0 = \overline{\gamma L}_1(\kappa) > 1$, where $\overline{\gamma L}_1(\kappa)$ is given explicitly in Theorem 3.2.4. Further on, $s_k, s_+ \in (1, \frac{2}{1+[\kappa]_+})$, for $k \geq 0$. By definition of s_+:

$$0 = \frac{s_k}{2 - s_k(1+\kappa)} - \frac{s_+}{2 - s_+(1+\kappa)} + \frac{s_+(2-s_+)(2-\kappa s_+)}{2 - s_+(1+\kappa)}.$$

Since $\frac{s_+(2-s_+)(2-\kappa s_+)}{2-s_+(1+\kappa)} \geq 0$, it implies:

$$\frac{s_+}{2 - s_+(1+\kappa)} \leq \frac{s_k}{2 - s_k(1+\kappa)} \iff s_k \leq s_+.$$

After simplifications in (3.15), s_+ are the roots of the function:

$$\begin{aligned} h(s_+) \coloneqq & -\kappa[2 - s_k(1+\kappa)]s_+^3 + 2(1+\kappa)[2 - s_k(1+\kappa)]s_+^2 + \\ & + 2[-3 + s_k(1+\kappa)]s_+ - 2s_k. \end{aligned}$$

The demonstration of the existence of a unique root s_k^+ is divided into scenarios of weak and strong convexity.

(a) **Case** $\kappa \in (-\infty, 0]$**.** Then $s_k \in [1, 2)$, $s_0 = \overline{\gamma L}_1(0) = \frac{3}{2}$ and the following holds:

i. $h(1) = -(1-\kappa)(2 - \kappa s_k) < 0$;
ii. $h(3/2) = \frac{-s_k(1-3\kappa)}{2} + \frac{9}{8}[2 - s_k(1+\kappa)] < 0$;
iii. $h(2) = 2(2 - s_k)$;
iv. $\frac{dh}{ds_+}(1) = (1+\kappa)(2 - \kappa s_k) \geq 0$ for $\kappa \in [-1, 1)$;
v. $\frac{dh}{ds_+}(\frac{3}{2}) = 2 + [2 - s_k(1+\kappa)](2 - \frac{3}{4}\kappa) > 0$ for $\kappa \in (-\infty, 0]$;
vi. $\frac{dh}{ds_+}(2) = 2 + 4(1-\kappa)[2 - s_k(1+\kappa)] > 0$ for $\kappa \in (-\infty, 0]$;
vii. $\frac{d^2 h(s_+)}{ds_+^2} = 2[2 - s_k(1+\kappa)][2 - \kappa - 3\kappa(s_+ - 1)] > 0$.

Therefore, h is strictly increasing and possesses a unique root $s_+ \in [\frac{3}{2}, 2)$.

(b) **Strongly convex case.** We restrict to $\kappa \in [0, 1)$ and $s_k \in [1, \frac{2}{1+\kappa})$. One can check the following statements:

i. $h(1) = -(1-\kappa)(2 - \kappa s_k) < 0$;
ii. $h(\frac{2}{1+\kappa}) = \frac{2(1-\kappa)^2[2 - s_k(1+\kappa)]}{(1+\kappa)^3} > 0$;
iii. $\frac{dh}{ds_+}(1) = (1+\kappa)(2 - \kappa s_k) > 0$;
iv. $\frac{dh}{ds_+}(\frac{2}{1+\kappa}) = 2(1+\kappa^2)[2 - s_k(1+\kappa)] + s_k\kappa(1+\kappa) + (1-\kappa)^2 > 0$;
v. $\frac{d^2 h(s_+)}{ds_+^2} = 2[2 - s_k(1+\kappa)]\left[\frac{3\kappa[2 - s_k^+(1+\kappa)] + 2[\kappa + (1-\kappa)^2]}{1+\kappa}\right] > 0$.

Therefore, h is strictly increasing and possesses a unique root $s_+ \in [1, \frac{2}{1+\kappa})$.

(ii) Since $s_{k+1} = s_+$ is uniquely defined and $s_k \leq s_+$, the sequence is monotonically increasing.

(iii) Due to its monotonicity, the sequence $\{s_k(\kappa)\}_{k=-1}^{\infty}$ has a limit, denoted by s_∞. We prove by contradiction that $s_\infty = \frac{2}{1+[\kappa]_+}$. Assume the contrary, i.e., $s_\infty \neq \frac{2}{1+[\kappa]_+}$, which is the upper bound of the interval of roots, so that $s_\infty \in [1, \frac{2}{1+[\kappa]_+})$. By replacing it in the definition of (3.15), it holds $\frac{s_\infty(2-s_\infty)(2-\kappa s_\infty)}{2-s_\infty(1+\kappa)} = 0$, which implies one of the solutions $s_\infty \in \{0, 2, \frac{2}{\kappa}\}$, which are all outside of the interval $[1, \frac{2}{1+[\kappa]_+})$, hence the contradiction!

(iv) When $\kappa = 0$, we define $c_j := 2 - s_j(0)$, $\forall j = \{-1, 0, \dots\}$. From the recurrence it results $2 - c_{j+1} - \frac{1}{c_{j+1}} + \frac{1}{c_j} = 0$. Since $\{s_j\}$ is monotonically increasing, with $s_0 = \frac{3}{2}$, we have that $c_0 = \frac{1}{2}$, $\{c_j\}$ is a decreasing sequence with $c_j \in (0, \frac{1}{2}]$, $\forall j \geq 0$, and $c_\infty = 0$. Thus, $\frac{1}{c_{j+1}} \geq \frac{1}{c_j} + \frac{3}{2}$, $\forall j \geq 0$. Telescoping for $j = 0, \dots, k-2$ we get $\frac{1}{c_{k-1}} \geq \frac{1}{c_0} + \frac{3}{2}(k-1)$, hence $c_{k-1} \leq \frac{2}{3k+1}$ and $s_{k-1} \geq 2 - \frac{1}{\frac{3}{2}k}$.

(v) When $\kappa \in (0, 1)$, we define $c_j := \frac{2}{1+\kappa} - s_j(\kappa)$, $\forall j = \{-1, 0, \dots\}$. Since $\{s_j\}$ is monotonically increasing, then $\{c_j\}$ is a monotonically decreasing sequence with $c_\infty = 0$. Using $s_0 = \overline{\gamma L}_1(\kappa) = \frac{3}{1+\kappa+\sqrt{1-\kappa+\kappa^2}}$, one can check that $c_0 \leq \min\{\frac{2}{1+\kappa}(\frac{1-\kappa}{1+\kappa})^2, \frac{1}{2}\}$, thus $c_j \in (0, \frac{1}{2}]$, $\forall j \geq 0$. We prove $\frac{c_{j+1}}{c_j} \leq (\frac{1-\kappa}{1+\kappa})^2$, $\forall j \geq 0$. By replacing in the recurrence of $s_j(\kappa)$ we get

$$\frac{\left(\frac{2}{1+\kappa} - c_{j+1}\right)\left(\frac{2\kappa}{1+\kappa} + c_{j+1}\right)\left(\frac{2}{1+\kappa} + \kappa c_{j+1}\right)}{c_{j+1}(1+\kappa)} - \frac{2}{c_{j+1}(1+\kappa)^2} + \frac{2}{c_j(1+\kappa)^2} = 0.$$

After algebraic manipulations, this is equivalent with

$$\begin{aligned}\frac{c_{j+1}}{c_j} &= 1 - \frac{1+\kappa}{2}\Big(\frac{2}{1+\kappa} - c_{j+1}\Big)\Big(\frac{2\kappa}{1+\kappa} + c_{j+1}\Big)\Big(\frac{2}{1+\kappa} + \kappa c_{j+1}\Big) \\ &= \Big(\frac{1-\kappa}{1+\kappa}\Big)^2 + \frac{\kappa(1+\kappa)c_{j+1}}{2}\Big[c_{j+1}^2 + 2c_{j+1}\frac{1+\kappa^3}{\kappa(1+\kappa)^2} - 4\frac{1+\kappa^3}{\kappa(1+\kappa)^3}\Big].\end{aligned}$$

For the expression within the square bracket, by exploiting its monotonicity with respect to c_{j+1} and evaluating it in $\frac{1}{2}$ we have that it is negative for any $c_{j+1} \in [0, \frac{1}{2}]$ and $\kappa \in (0, 1)$. Therefore, we get $\frac{c_{j+1}}{c_j} \leq (\frac{1-\kappa}{1+\kappa})^2$. By telescoping for $j = 0, \dots, k-2$, this implies $c_{k-1} \leq c_0(\frac{1-\kappa}{1+\kappa})^{2(k-1)} \leq (\frac{1-\kappa}{1+\kappa})^{2k}$, hence $s_{k-1} \geq \frac{2}{1+\kappa} - (\frac{1-\kappa}{1+\kappa})^{2k}$.

□

3.A.4 Performance bounds for dynamic stepsizes

In this section we give the proof of Theorem 3.3.1, within which we also show the results from Corollary 3.3.2 (for strongly convex function), Corollary 3.3.1 (for convex functions) and Corollary 3.3.3 (for weakly convex functions).

Proof of Theorem 3.3.1. By telescoping (N4SD) with stepsizes $\gamma_i L$:

$$f_0 - f_N \geq \frac{\gamma_{N-1}L}{2-\gamma_{N-1}L-\gamma_{N-1}\mu}\frac{\|g_N\|^2}{2L} + \frac{\gamma_0 L[(2-\gamma_0 L)(2-\gamma_0\mu)-1]}{2-\gamma_0 L-\gamma_0\mu}\frac{\|g_0\|^2}{2L} + \\ + \sum_{i=1}^{N-1}\frac{\|g_i\|^2}{2L}\Big[\frac{\gamma_{i-1}L}{2-\gamma_{i-1}L-\gamma_{i-1}\mu} + \frac{\gamma_i L[(2-\gamma_i L)(2-\gamma_i\mu)-1]}{2-\gamma_i L-\gamma_i\mu}\Big]. \tag{3.45}$$

Given that all weights are nonnegative (the fact which we prove below), the bound (3.16) is obtained by taking the minimum gradient norm:

$$f_0 - f_N \geq \min_{0\leq i\leq N}\left\{\frac{\|g_i\|^2}{2L}\right\}\sum_{i=0}^{N-1}\frac{\gamma_i L(2-\gamma_i L)(2-\gamma_i\mu)}{2-\gamma_i L-\gamma_i\mu}.$$

For $\kappa < 0$, we define $\tilde{N} := \arg\max_{0\leq i\leq N-1}\{s_i \leq (\gamma L)^*(\kappa)\}$.
Case (i): $\gamma_i L = s_i(\kappa)$ for $i = 0, \dots, N-1$. This corresponds to the (strongly) convex case ($\kappa \in [0,1)$) or weakly convex ($\kappa < 0$) with $N \leq \tilde{N}$. By the definition of $s_i(\kappa)$, all weights but the one of $\|g_N\|$ vanish and the inequality (3.45) reduces to:

$$f_0 - f_N \geq \frac{s_{N-1}}{2-s_{N-1}(1+\kappa)}\frac{\|g_N\|^2}{2L}, \tag{3.46}$$

which is of type (♣) and leads to the rate (3.18) from Corollary 3.3.2.
Moreover, in the convex case the expression of $s_k(0)$ can be computed analytically from (3.15) as the root belonging to $[1,2)$ of:

$$\frac{s_k(0)(3-2s_k(0))}{2-s_k(0)} + \frac{s_{k-1}(0)}{2-s_{k-1}(0)} = 0.$$

This solution is $s_k(0) = \frac{3-2s_{k-1}(0)+\sqrt{9-4s_{k-1}(0)}}{2(2-s_{k-1}(0))}$, which equivalently rewrites as in the hypothesis of Corollary 3.3.1.

Case (ii): $\gamma_i L = \min\{s_i(\kappa), (\gamma L)^*(\kappa)\}$ for $i = 0, \dots, N-1$. This case corresponds to $\kappa < 0$ and $N > \tilde{N}$. Then inequality (3.45) becomes

$$f_0 - f_N \geq \left[\frac{s_{\tilde{N}}}{2 - s_{\tilde{N}}(1+\kappa)} + \frac{(\gamma L)^*[(2-(\gamma L)^*)(2-\kappa(\gamma L)^*)-1]}{2-(\gamma L)^*(1+\kappa)}\right]\frac{\|g_{\tilde{N}+1}\|^2}{2L} + \sum_{i=\tilde{N}+2}^{N-1}\left[\frac{(\gamma L)^*(2-(\gamma L)^*)(2-\kappa(\gamma L)^*)}{2-(\gamma L)^*(1+\kappa)}\frac{\|g_i\|^2}{2L}\right] + \frac{(\gamma L)^*}{2-(\gamma L)^*(1+\kappa)}\frac{\|g_N\|^2}{2L}. \tag{3.47}$$

It remains to prove the nonnegativity of $\|g_{\tilde{N}+1}\|^2$'s weight, namely:

$$\frac{s_{\tilde{N}}}{2 - s_{\tilde{N}}(1+\kappa)} + \frac{(\gamma L)^*[(2-(\gamma L)^*)(2-\kappa(\gamma L)^*)-1]}{2-(\gamma L)^*(1+\kappa)} \geq 0.$$

By definition of (3.15), it holds:

$$\frac{s_{\tilde{N}}}{2 - s_{\tilde{N}}(1+\kappa)} = -\frac{s_{\tilde{N}+1}[(2-s_{\tilde{N}+1})(2-\kappa s_{\tilde{N}+1})-1]}{2-s_{\tilde{N}+1}(1+\kappa)},$$

with $s_{\tilde{N}+1} > (\gamma L)^*(\kappa)$. The inequality to show is equivalent to

$$-\frac{s_{\tilde{N}+1}[(2-s_{\tilde{N}+1})(2-\kappa s_{\tilde{N}+1})-1]}{2-s_{\tilde{N}+1}(1+\kappa)} \geq -\frac{(\gamma L)^*[(2-(\gamma L)^*)(2-\kappa(\gamma L)^*)-1]}{2-(\gamma L)^*(1+\kappa)},$$

which is implied by the monotone decreasing of the function $\sigma\colon [1,2) \to \mathbb{R}$, $\sigma(t) := \frac{t[(2-t)(2-\kappa t)-1]}{2-t(1+\kappa)}$, having $\frac{d\sigma(t)}{dt} = -\frac{2(1-\kappa t)(t-1)[1+2-t(1+\kappa)]}{[2-t(1+\kappa)]^2} \leq 0$.
Taking the minimum gradient norm in (3.47), with $p(\gamma L, \gamma\mu)$ defined in (3.7):

$$f_0 - f_N \geq \min_{\tilde{N}+1 \leq i \leq N}\left\{\frac{\|g_i\|^2}{2L}\right\}\left[\frac{s_{\tilde{N}}}{2-s_{\tilde{N}}(1+\kappa)} + p((\gamma L)^*, \kappa(\gamma L)^*)(N-\tilde{N}-1)\right]. \tag{3.48}$$

Since the sequence $\{s_k\}$ is increasing, the inequality $\frac{s_{\tilde{N}}}{2-s_{\tilde{N}}(1+\kappa)} \geq \frac{s_{N-1}}{2-s_{N-1}(1+\kappa)}$ is equivalent to $\tilde{N} \geq N-1$. Consequently, the merged r.h.s. of (3.46) and (3.48) can be expressed as a maximum:

$$\max\left\{\frac{s_{N-1}}{2-s_{N-1}(1+\kappa)}, \frac{s_{\tilde{N}}}{2-s_{\tilde{N}}(1+\kappa)} + p((\gamma L)^*, \kappa(\gamma L)^*)(N-\tilde{N}-1)\right\} = p((\gamma L)^*, \kappa(\gamma L)^*)N + \max_{0\leq k\leq N}\left\{\frac{s_{k-1}}{2-s_{k-1}(1+\kappa)} - p((\gamma L)^*, \kappa(\gamma L)^*)k\right\},$$

which is exactly as in (3.19) from Corollary 3.3.3. □

3.B Proof of optimal point's characterization for smooth functions

Proof of Lemma 3.6.1. We extend the [43, Theorem 7] to smooth functions with any lower curvature μ, generalizing from the original convex ($\mu = 0$) and non-necessarily ($\mu = -L$) cases. Consider the function

$$Z(y) \coloneqq \min_{\alpha \in \Delta_\mathcal{I}} \Big\{ \tfrac{L-\mu}{2} \Big\| y - \sum_{i \in \mathcal{I}} \alpha_i \big[x_i - \tfrac{1}{L-\mu} \big(g_i - \mu x_i \big) \big] \Big\|^2 + \sum_{i \in \mathcal{I}} \alpha_i \big(f_i - \tfrac{\mu}{2} \|x_i\|^2 - \tfrac{1}{2(L-\mu)} \|g_i - \mu x_i\|^2 \big) \Big\}$$

where $\Delta_\mathcal{I}$ is the $|\mathcal{I}|$-dimensional unit simplex:

$$\Delta_\mathcal{I} \coloneqq \big\{ \alpha \in \mathbb{R}^n : \sum_{i \in \mathcal{I}} \alpha_i = 1, \alpha_i \geq 0, \forall i \in \mathcal{I} \big\}.$$

The function Z is the primal interpolation function $W_\mathcal{T}^C$ as defined in [42, Definition 2.1]. One can check this by replacing in the definition $C \leftarrow \{0\}$, $L \leftarrow L - \mu$ and

$$\mathcal{T} \leftarrow \big\{ \big(x_i, g_i - \mu x_i, f_i - \tfrac{\mu}{2} \|x_i\|^2 \big) \big\}_{i \in \mathcal{I}}.$$

The set $\mathcal{T}$ is $\mathcal{F}_{0,L-\mu}$-interpolable. Hence, from [42, Theorem 1] it follows that Z is convex, with upper curvature $L - \mu$, and satisfies

$$Z(x_i) = f_i - \frac{\mu}{2} \|x_i\|^2 \quad \text{and} \quad \nabla Z(x_i) = g_i - \mu x_i.$$

Consider the function

$$\hat{W}(y) \coloneqq Z(y) + \tfrac{\mu}{2} \|y\|^2,$$

which by curvature shifting belongs to the function class $\mathcal{F}_{\mu,L}$ and satisfies

$$\hat{W}(x_i) = f_i \quad \text{and} \quad \nabla \hat{W}(x_i) = g_i.$$

Using algebraic manipulations, $\hat{W}$ can be expressed as

$$\hat{W}(y) = \min_{\alpha \in \Delta_\mathcal{I}} \Big\{ \tfrac{L}{2} \Big\| y - \sum_{i \in \mathcal{I}} \alpha_i \big(x_i - \tfrac{1}{L} g_i \big) \Big\|^2 + \tfrac{L}{2} \tfrac{\mu}{L-\mu} \Big\| \sum_{i \in \mathcal{I}} \alpha_i \big(x_i - \tfrac{1}{L} g_i \big) \Big\|^2 + \sum_{i \in \mathcal{I}} \alpha_i \big(f_i - \tfrac{1}{2L} \|g_i\|^2 - \tfrac{L}{2} \tfrac{\mu}{L-\mu} \Big\| x_i - \tfrac{1}{L} g_i \Big\|^2 \big) \Big\}. \tag{3.49}$$

Further, we *lower bound* the first squared norm by 0:

$$\hat{W}(y) \geq \min_{\alpha \in \Delta_{\mathcal{I}}} \Big\{ \tfrac{L}{2} \tfrac{\mu}{L-\mu} \Big\| \sum_{i \in \mathcal{I}} \alpha_i \big(x_i - \tfrac{1}{L} g_i\big) \Big\|^2 +$$

$$\sum_{i \in \mathcal{I}} \alpha_i \big(f_i - \tfrac{1}{2L} \|g_i\|^2 - \tfrac{L}{2} \tfrac{\mu}{L-\mu} \big\|x_i - \tfrac{1}{L} g_i\big\|^2\big) \Big\}.$$

Applying Jensen inequality for the squared norm, which is a convex function, we get:

$$\Big\| \sum_{i \in \mathcal{I}} \alpha_i \big(x_i - \tfrac{1}{L} g_i\big) \Big\|^2 \leq \sum_{i \in \mathcal{I}} \alpha_i \big\|x_i - \tfrac{1}{L} g_i\big\|^2.$$

Then for $\mu \leq 0$ we can further lower bound the above inequality and obtain:

$$\begin{aligned} \hat{W}(y) &\geq \min_{\alpha \in \Delta_{\mathcal{I}}} \Big\{ \sum_{i \in \mathcal{I}} \alpha_i \big(f_i - \tfrac{1}{2L} \|g_i\|^2\big) \Big\} \\ &= \min_{i \in \mathcal{I}} \;\; \big\{ f_i - \tfrac{1}{2L} \|g_i\|^2 \big\} = f_{i_*} - \tfrac{1}{2L} \|g_{i_*}\|^2. \end{aligned}$$

For the *upper bound*, in (3.49) we take $y := x_{i_*} - \frac{1}{L} g_{i_*}$ and $\alpha = e_{i_*}$ (the i_*-th unit vector):

$$\begin{aligned} \hat{W}\big(x_{i_*} - \tfrac{1}{L} g_{i_*}\big) &\leq \tfrac{L}{2} \tfrac{\mu}{L-\mu} \big\|x_{i_*} - \tfrac{1}{L} g_{i_*}\big\|^2 + f_{i_*} - \tfrac{1}{2L} \|g_{i_*}\|^2 - \tfrac{L}{2} \tfrac{\mu}{L-\mu} \big\|x_{i_*} - \tfrac{1}{L} g_{i_*}\big\|^2 \\ &= f_{i_*} - \tfrac{1}{2L} \|g_{i_*}\|^2. \end{aligned}$$

Therefore, from both lower and upper bounds it follows $\hat{W}\big(x_{i_*} - \frac{1}{L} g_{i_*}\big) = f_{i_*} - \frac{1}{2L} \|g_{i_*}\|^2$. □

3.C Tightness proofs (details)

3.C.1 Sublinear regimes with stepsizes $\gamma_i L \in (1, \overline{\gamma L}_1(\kappa)]$

Proof of Proposition 3.6.5. The equality case from Theorem 3.2.4's proof implies all gradient norms equal with each other as defined in (3.38). By replacing this expression in the equality conditions from Lemma 3.5.2, identity (3.39) results. Identity (3.40) is implied by the equality case in (N4SD):

$$f_i - f_{i+1} = \frac{\gamma_i L (2 - \gamma_i \mu)(2 - \gamma_i L)}{2 - \gamma_i L - \gamma_i \mu} \frac{U^2}{2L}.$$

The normalized inner product after one iteration is $c(\gamma_i L, \gamma_i \mu) \in (-1, 1)$, therefore $\theta(\gamma L, \gamma\mu)$ is well defined in (3.42). Moreover, recurrence (3.41) satisfies the necessary condition (3.39) since

$$\langle g_i, g_{i+1}\rangle = U^2 \cos\big((-1)^i \theta(\gamma_i L, \gamma_i \mu)\big) = U^2\, c(\gamma_i L, \gamma_i \mu).$$

(3.43) holds because all gradient norms are equal and the function values decrease. □

Proof of Proposition 3.6.6 (cont.) We check that all interpolation conditions ($Q_{[i,j]}$) and ($Q_{[j,i]}$) hold $\forall (i,j) \in \{0, 1, \dots, N\}$: (without loss of generality $j > i$)

$$Q_{[i,j]}\colon f_i - f_j - \gamma\Big\langle g_j\,,\sum_{k=i}^{j-1} g_k\Big\rangle \geq$$

$$\frac{1}{2(L-\mu)}\Big(\big\|g_i - g_j\big\|^2 + \gamma L\,\gamma\mu\Big\|\sum_{k=i}^{j-1} g_k\Big\|^2 - 2\gamma\mu\Big\langle g_i - g_j, \sum_{k=i}^{j-1} g_k\Big\rangle\Big);$$

$$Q_{[j,i]}\colon f_j - f_i + \gamma\Big\langle g_i\,,\sum_{k=i}^{j-1} g_k\Big\rangle \geq$$

$$\frac{1}{2(L-\mu)}\Big(\big\|g_i - g_j\big\|^2 + \gamma L\,\gamma\mu\Big\|\sum_{k=i}^{j-1} g_k\Big\|^2 - 2\gamma\mu\Big\langle g_i - g_j, \sum_{k=i}^{j-1} g_k\Big\rangle\Big).$$

Note the same expressions of the r.h.s. in both inequalities. Let $\Delta := f_0 - f_N$ and recall $c(\gamma L, \gamma\mu) = \frac{1+(1-\gamma L)(1-\gamma\mu)}{2-\gamma L-\gamma\mu}$. Then $1 + c = \frac{(2-\gamma L)(2-\gamma\mu)}{2-\gamma L-\gamma\mu}$ and

$$U^2 := \frac{\Delta}{\frac{1}{2L}\frac{\gamma L(2-\gamma L)(2-\gamma\mu)}{2-\gamma(L+\mu)} N} = \frac{2L\Delta}{(1+c)N}. \tag{3.50}$$

- **Case $j - i$ even.** Then $g_i = g_j$ and we use the identity:

$$\Big\langle g_i\,, \sum_{k=i}^{j-1} g_k\Big\rangle = \frac{U^2(1+c)(j-i)}{2}$$

to rewrite the interpolation inequalities as follows:

$$Q_{[i,j]}: \quad \frac{\Delta(j-i)}{N} - \frac{\gamma U^2(1+c)(j-i)}{2} \geq \frac{\gamma L\gamma\mu}{2(L-\mu)}\Big\|\sum_{k=i}^{j-1} g_k\Big\|^2;$$

$$Q_{[j,i]}: \quad \frac{\Delta(i-j)}{N} + \frac{\gamma U^2(1+c)(j-i)}{2} \geq \frac{\gamma L\gamma\mu}{2(L-\mu)}\Big\|\sum_{k=i}^{j-1} g_k\Big\|^2.$$

By replacing the definition of U^2, the l.h.s. in both inequalities equals zero, while the r.h.s. is non-positive because $\mu \leq 0$; therefore, both interpolation inequalities are valid for $i + j$ even.

• **Case** $j - i$ **odd.** Then $g_i \neq g_j$ and $\langle g_i, g_j \rangle = cU^2$, and we use the following identities:

$$\Big\langle g_i\,, \sum_{k=i}^{j-1} g_k \Big\rangle = \frac{U^2(j-i)(1+c)}{2} + \frac{U^2(1-c)}{2};$$

$$\Big\langle g_j\,, \sum_{k=i}^{j-1} g_k \Big\rangle = \frac{U^2(j-i)(1+c)}{2} - \frac{U^2(1-c)}{2};$$

$$\Big\langle g_i - g_j, \sum_{k=i}^{j-1} g_k \Big\rangle = U^2(1-c);$$

$$\Big\| \sum_{k=i}^{j-1} g_k \Big\|^2 = \frac{U^2}{2}\big[(1+c)(j-i)^2 + (1-c)\big];$$

$$\big\| g_i - g_j \big\|^2 = 2(1-c)U^2.$$

After substitutions, the r.h.s. and l.h.s. of both inequalities $Q_{[i,j]}$ and $Q_{[j,i]}$ are the same:

$$\text{r.h.s.}_{Q_{[i,j]}} = \text{r.h.s.}_{Q_{[j,i]}} =$$

$$\frac{U^2}{4(L-\mu)}\big[\gamma L\gamma\mu(1+c)(j-i)^2 + (1-c)\big(\gamma L\gamma\mu - 4\gamma\mu + 4\big)\big]$$

and, respectively:

$$\text{l.h.s.}_{Q_{[i,j]}}: \quad \frac{\Delta(j-i)}{N} - \gamma\Big\langle g_j, \sum_{k=i}^{j-1} g_k \Big\rangle = \frac{\gamma U^2(1-c)}{2};$$

$$\text{l.h.s.}_{Q_{[j,i]}}: \quad \frac{\Delta(i-j)}{N} + \gamma\Big\langle g_i, \sum_{k=i}^{j-1} g_k \Big\rangle = \frac{\gamma U^2(1-c)}{2}.$$

It remains to prove (with $j - i = 1, 3, \dots$):

$$\frac{\gamma U^2(1-c)}{2} \geq \frac{U^2}{4(L-\mu)}\Big[\gamma L\,\gamma\mu(1+c)(j-i)^2 + (1-c)\big(\gamma L\,\gamma\mu - 4\gamma\mu + 4\big)\Big],$$

equivalent to

$$0 \geq \gamma L\gamma\mu(1+c)(j-i)^2 + (1-c)\big[\gamma L\gamma\mu - 2\gamma(L+\mu) + 4\big].$$

By using identity

$$(1-c)\big[\gamma L\gamma\mu - 2\gamma(L+\mu)+4\big] = -\gamma L\gamma\mu(1+c),$$

the inequality to prove becomes:

$$0 \geq \gamma L\gamma\mu(1+c)\big[(j-i)^2-1\big].$$

This is true since $j-i\geq 1$, $\mu\leq 0$ and $c(\gamma L,\gamma\mu)\in[-1,1]$; the equality case holds for $j=i+1$, i.e., the necessary conditions from equality in the distance-1 interpolation inequalities used to derive the upper bound.

We have shown that the proposed set of triplets satisfy the interpolation conditions for all pairs (i,j), hence, following Theorem 2.2.1, it is $\mathcal{F}_{\mu,L}$-interpolable. □

3.C.2 Proof of the three-dimensional worst-case example from Proposition 3.6.8

Preliminaries. We denote $\eta := 1-\gamma\mu \in (1,\infty)$, $\rho := 1-\gamma L\in(-1,0)$, $\tilde{U} := \sqrt{\frac{(1-\rho^2)(\eta^2-1)}{\eta^2-\rho^2}}$, and rewrite the expressions defining the triplets as:

$$g_i = U\tilde{U}\begin{bmatrix} \frac{\rho^{i-\bar{N}}}{\sqrt{1-\rho^2}} \\ \frac{\eta^{i-\bar{N}}}{\sqrt{\eta^2-1}} \\ 0 \end{bmatrix} \qquad \forall i = 0,1,\ldots,\bar{N}+1;$$

$$g_{i+1} = \rho g_i + (\eta-\rho)\cos(\theta_{i+1})\begin{bmatrix} 0 & \\ & R(\theta_{i+1}) \end{bmatrix} g_i \quad \forall i=\bar{N},\ldots,N. \tag{3.51}$$

$$f_i = \Delta\frac{N-i+\frac{(1-\rho)(1-\eta)}{\eta-\rho}\sum_{j=1}^{\bar{N}-i}T_j(\rho,\eta)}{N+\frac{(1-\rho)(1-\eta)}{\eta-\rho}\sum_{j=1}^{\bar{N}}T_j(\rho,\eta)} \quad i=0,1,\ldots,N. \tag{3.52}$$

The recurrence (3.51) also works to define $g_{\bar{N}+1}$ due to $q_{\bar{N}+1}=0$ and $\theta_{\bar{N}+1}=0$. We consider

$$S_k := \sum_{p=0}^{k-(\bar{N}+2)} \rho^{2p} = \frac{1-\rho^{2(k-(\bar{N}+1))}}{1-\rho^2}, \tag{3.53}$$

satisfying the properties $S_{k+1} = 1+\rho^2 S_k$ and $S_\infty := \lim_{k\to\infty} S_k = \frac{1}{1-\rho^2}$. For each $k=\bar{N}+1,\bar{N}+2,\ldots,N$ we have that

$$q_k = \tan(\theta_k) = (-1)^{k-(\bar{N}+1)}\sqrt{(\eta^2-1)S_k}.$$

1. Proof of reaching the performance bound. We demonstrate that $\min_{0 \le i \le \bar{N}} \|g_i\| = U$.
Note that:

$$\|g_i\|^2 = U^2 \tilde{U}^2 \big(\tfrac{\eta^{2(i-\bar{N})}}{\eta^2-1} + \tfrac{\rho^{2(i-\bar{N})}}{1-\rho^2}\big) \quad i = 0, \dots, \bar{N}+1; \tag{3.54}$$

in particular, $\|g_{\bar{N}}\|^2 = \|g_{\bar{N}+1}\|^2 = U^2$. For the latter, one can use the identity $\frac{\eta^2}{\eta^2-1} - \frac{\rho^2}{1-\rho^2} = \frac{1}{\eta^2-1} - \frac{1}{1-\rho^2}$, obtained by adding 1 to both ratios. We show that all gradients with indices $i \ge \bar{N}$ have norm U, namely $\|g_i\|^2 = U^2$ for any $i = \bar{N}, \dots, N$.
Recurrence (3.51) implies that the first component of g_i is always $U\tilde{U}\frac{\rho^{i-\bar{N}}}{\sqrt{1-\rho^2}}$, being a scaling with ρ of the first component of g_{i-1}. Let $v_i := \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} g_i$ be the vector of the last two entries in g_i. Thus showing that $\|g_i\|^2 = U^2$ for $i \ge \bar{N}$ reduces to prove that

$$\|v_i\|^2 = U^2 \Big(1 - \tilde{U}^2 \frac{\rho^{2(i-\bar{N})}}{1-\rho^2}\Big), \ \forall i = \bar{N}, \bar{N}+1, \dots, N. \tag{3.55}$$

We show this by induction. Firstly, it rewrites as

$$\|v_i\|^2 = U^2 \tilde{U}^2 \frac{1+q_{i+1}^2}{\eta^2-1}, \ \forall i = \bar{N}, \bar{N}+1, \dots, N. \tag{3.56}$$

Cases $i = \{\bar{N}, \bar{N}+1\}$ result from the expressions of $g_{\bar{N}}$ and $g_{\bar{N}+1}$. Assuming that (3.55) holds for some index k, we show its validity for $k+1$. The recurrence of the gradients implies

$$v_{k+1} = \rho v_k + (\eta - \rho) \cos(\theta_{k+1}) R(\theta_{k+1}) v_k.$$

Taking the norm, expanding the squares, using the properties that the rotation matrix preserves the vector's norm and $\langle v_k, R(\theta_{k+1}) v_k \rangle = \|v_k\|^2 \cos(\theta_{k+1})$ yields

$$\|v_{k+1}\|^2 = \|v_k\|^2 \big[\rho^2 + (\eta^2 - \rho^2) \cos^2(\theta_{k+1})\big].$$

Using $\cos^2(\theta_{k+1}) = \frac{1}{1+q_{k+1}^2}$ and the identity $\eta^2 + \rho^2 q_{k+1}^2 = 1 + q_{k+2}^2$ we get

$$\|v_{k+1}\|^2 = \|v_k\|^2 \frac{1+q_{k+2}^2}{1+q_{k+1}^2},$$

where replacing $\|v_k\|$ using (3.56) yields the expression in (3.55) written for $k+1$.

Since $U^2 = \frac{2L\,\Delta}{\gamma L P_N(\gamma L,\gamma\mu)}$ and $f_0 - f_N = \Delta$, the triplets reach the performance bound with respect to $f(x_0) - f(x_N)$ from Proposition 3.6.8. To prove the bound with respect to $f(x_0) - f_*$, it remains to show that the required assumption in Lemma 3.6.2 holds, namely

$$N \in \operatorname*{arg\,min}_{0\le i\le N} \left\{ f_i - \tfrac{1}{2L}\|g_i\|^2 \right\}. \tag{3.57}$$

After replacing the expressions of gradient norms in (3.52) and several simplifications we get that, for all $i = 0, \dots, N$, it holds:

$$f_i - \frac{\|g_i\|^2}{2L} = f_N - \frac{U^2}{2L} + \frac{U^2}{2L} \frac{(1-\rho)(1+\rho)(1+\eta)}{\eta+\rho} \Big[N - \bar{N} + \sum_{j=1}^{\bar{N}-i} \eta^{-2j} \Big],$$

where, with an abuse of notation, the sum is zero for $i \ge \bar{N}$. Its minimum is reached for $i = N$ and therefore identity (3.57) holds.

2. The set of triplets $\mathcal{T}_{\mathcal{I}}$ is $\mathcal{F}_{\mu,L}$-interpolable. We show that all interpolation conditions (2.7) are satisfied, namely for any i,j with $0 \le i,j \le N$, $i \ne j$, it holds that $Q_{i,j} \ge 0$, where

$$\begin{aligned} Q_{i,j} &\coloneqq f_i - f_j - \langle g_j, x_i - x_j \rangle \\ &\qquad - \frac{1}{2L}\|g_i - g_j\|^2 - \frac{\mu}{2L(L-\mu)} \|g_i - g_j - L(x_i - x_j)\|^2. \end{aligned}$$

We divide the analysis into the cases: $i > j$ and $i < j$, showing that $Q_{i,j}$ rewrites as a sum of nonnegative contributions in each of them. The central challenge lies in finding an exact expression for the inner product $\langle g_{k+1} - \eta g_k\,,\, g_{i+1} - \rho g_i \rangle$ for all $i, k \ge \bar{N}$, which we obtain below in (3.61).

Preliminaries. For simplicity of notation and without loss of generality, we consider $U = 1$ since all gradients are multiplied by $U = \min_{0\le i\le N} \|\nabla f(x_i)\|$.

Following (3.53), observe that $q_{\bar{N}+1} = 0$, $q_{\bar{N}+2} = -\sqrt{\eta^2 - 1}$ and $|q_k|$ is monotonically increasing towards $|q_\infty| = \sqrt{\frac{\eta^2-1}{1-\rho^2}}$. Consequently, $|\theta_k| = \arctan(|q_k|) \in (0, \frac{\pi}{2})$, $|\theta_k|$ monotonically increases towards $\arctan(|q_\infty|)$ and $\operatorname{sgn}(\theta_k) = (-1)^{k-(\bar{N}+1)}$.

For all $k = 0, \dots, \bar{N}+1$, the gradients read as $g_k = \tilde{U}[\frac{\rho^{k-\bar{N}}}{\sqrt{1-\rho^2}}, \frac{\eta^{k-\bar{N}}}{\sqrt{\eta^2-1}}, 0]^\top$. Furthermore, for all $k \le \bar{N}$ and $i = 0, \dots, N-1$ it holds that:

$$\langle g_{k+1} - \eta g_k\,,\, g_{i+1} - \rho g_i \rangle = \langle [*, 0, 0]^\top\,,\, [0, *, *]^\top \rangle = 0, \tag{3.58}$$

where $*$ denotes for the entries of the three dimensional vectors multiplied by zero. Only the first component is $g_{k+1} - \eta g_k$ is nonzero (for $k \leq \bar{N}$), while the first component in $g_{i+1} - \rho g_i$ is always zero for any $i = 0, \ldots, N-1$.

Calculation of $\langle g_{k+1} - \eta g_k \, , \, g_{i+1} - \rho g_i \rangle$ for any $i, k \geq \bar{N}$. Identity (3.56) implies for all $k = \bar{N}, \bar{N}+1, \ldots, N-1$ that $\|v_k\| \cos(\theta_{k+1}) = \frac{\tilde{U}}{\sqrt{\eta^2-1}}$ and thus:

$$v_{k+1} - \rho v_k = (\eta - \rho) \frac{\tilde{U}}{\sqrt{\eta^2 - 1}} R(\theta_{k+1}) \frac{v_k}{\|v_k\|};$$

$$\begin{aligned} v_{k+1} - \eta v_k &= -(\eta - \rho)\big[I_2 - \cos(\theta_{k+1}) R(\theta_{k+1})\big] v_k \\ &\overset{(3.56)}{=} -(\eta - \rho) \frac{\tilde{U}}{\sqrt{\eta^2 - 1}} q_{k+1} R\big(-\frac{\pi}{2}\big) R(\theta_{k+1}) \frac{v_k}{\|v_k\|}. \end{aligned}$$

Then we have

$$\langle g_{k+1} - \eta g_k \, , \, g_{i+1} - \rho g_i \rangle = \frac{-(\eta - \rho)^2 \tilde{U}^2}{\eta^2 - 1} q_{k+1} \cos(\beta_{k,i}), \tag{3.59}$$

where $\beta_{k,i} := \angle\big(R(-\frac{\pi}{2}) R(\theta_{k+1}) \frac{v_k}{\|v_k\|} \, , \, R(\theta_{i+1}) \frac{v_i}{\|v_i\|}\big)$. We normalize all vectors to transform the problem to maneuvering angles on the unit circle. Let $\phi(y)$ be the oriented angle of a two-dimensional vector y. We express the angles as follows:

$$\beta_{k,i} = \phi\Big(R(\theta_{i+1}) \frac{v_i}{\|v_i\|}\Big) - \phi\Big(R(-\frac{\pi}{2}) R(\theta_{k+1})\Big);$$

$$-\frac{\pi}{2} = \phi\Big(R(-\frac{\pi}{2}) R(\theta_{k+1}) \frac{v_k}{\|v_k\|}\Big) - \phi\Big(R(\theta_{k+1}) \frac{v_k}{\|v_k\|}\Big);$$

$$\theta_{k+1} = \phi\Big(R(\theta_{k+1}) \frac{v_k}{\|v_k\|}\Big) - \phi\Big(\frac{v_k}{\|v_k\|}\Big);$$

$$\theta_{i+1} = \phi\Big(R(\theta_{i+1}) \frac{v_i}{\|v_i\|}\Big) - \phi\Big(\frac{v_i}{\|v_i\|}\Big).$$

Using these identities yields

$$\beta_{k,i} = \frac{\pi}{2} + \theta_{i+1} - \theta_{k+1} + \phi\Big(\frac{v_i}{\|v_i\|}\Big) - \phi\Big(\frac{v_k}{\|v_k\|}\Big).$$

For any $l \geq \bar{N}$, let $\delta_{l,l+1} := \phi\big(\frac{v_{l+1}}{\|v_{l+1}\|}\big) - \phi\big(\frac{v_l}{\|v_l\|}\big)$ be the oriented angle from $\frac{v_l}{\|v_l\|}$ to $\frac{v_{l+1}}{\|v_{l+1}\|}$ and $\epsilon_{l,l+1} := \theta_{l+2} - \theta_{l+1} + \delta_{l,l+1}$. We rewrite $\beta_{k,i}$ as

$$\beta_{k,i} = \frac{\pi}{2} + \sum_{l=\bar{N}}^{i-1} \epsilon_{l,l+1} - \sum_{l=\bar{N}}^{k-1} \epsilon_{l,l+1} = \begin{cases} \frac{\pi}{2} + \sum_{l=k}^{i-1} \epsilon_{l,l+1}, & \text{if } i \geq k; \\ \frac{\pi}{2} - \sum_{l=i}^{k-1} \epsilon_{l,l+1}, & \text{if } i < k. \end{cases}$$

We compute a closed form of $\sin(\epsilon_{l,l+1})$ with $l \geq \bar{N}$:

$$\sin(\epsilon_{l,l+1}) = \sin(\delta_{l,l+1})\cos(\theta_{l+2}-\theta_{l+1}) + \cos(\delta_{l,l+1})\sin(\theta_{l+2}-\theta_{l+1}).$$

Using that $q_l = \tan(\theta_l)$, we get:

$$\cos(\theta_{l+2}-\theta_{l+1}) = \frac{1+q_{l+1}q_{l+2}}{\sqrt{(1+q_{l+2}^2)(1+q_{l+1}^2)}};$$

$$\sin(\theta_{l+2}-\theta_{l+1}) = \frac{q_{l+2}-q_{l+1}}{\sqrt{(1+q_{l+2}^2)(1+q_{l+1}^2)}}.$$

Since $\cos(\delta_{l,l+1}) = \langle \frac{v_l}{\|v_l\|}, \frac{v_{l+1}}{\|v_{l+1}\|}\rangle$ and $\sin(\delta_{l,l+1}) = \frac{v_l}{\|v_l\|} \times \frac{v_{l+1}}{\|v_{l+1}\|}$, we obtain

$$\cos(\delta_{l,l+1}) = \frac{\eta+\rho q_{l+1}^2}{\sqrt{(1+q_{l+2}^2)(1+q_{l+1}^2)}};\ \sin(\delta_{l,l+1}) = \frac{(\eta-\rho)q_{l+1}}{\sqrt{(1+q_{l+2}^2)(1+q_{l+1}^2)}}.$$

Replacing in the expansion of $\sin(\epsilon_{l,l+1})$, after simplifications we get:

$$\sin(\epsilon_{l,l+1}) = \frac{\eta q_{l+2}-\rho q_{l+1}}{1+q_{l+2}^2}.$$

Substituting q_{l+1} and q_{l+2} and amplifying by the conjugate square root yields:

$$\sin(\epsilon_{l,l+1}) = (-1)^{l-(\bar{N}+1)}\frac{\sqrt{\eta^2-1}}{\eta\sqrt{S_{l+2}}-\rho\sqrt{S_{l+1}}}.$$

For simplicity of notation, we define the angle α_l as

$$\alpha_l := \arcsin\left(\frac{\sqrt{\eta^2-1}}{\eta\sqrt{S_{l+2}}-\rho\sqrt{S_{l+1}}}\right),\ \forall l \geq \bar{N}. \tag{3.60}$$

Since $\rho < 0$, $\eta > 1$ and S_l is monotonically increasing with l, we have that α_l is monotonically decreasing with l, hence $\alpha_l \in [\alpha_\infty, \alpha_{\bar{N}}] \subset (0, \frac{\pi}{2})$, where $\alpha_{\bar{N}} = \arcsin\left(\sqrt{1-\frac{1}{\eta^2}}\right)$ and $\alpha_\infty = \arcsin\left(\frac{\sqrt{(\eta^2-1)(1-\rho^2)}}{\eta-\rho}\right)$.
Using $q_{k+1} = \sqrt{(\eta^2-1)S_{k+1}}(-1)^{k-\bar{N}}$, after substituting everything into (3.59) we get

$$\langle g_{k+1}-\eta g_k\,,\, g_{i+1}-\rho g_i\rangle =$$

$$\frac{-(\eta-\rho)^2\tilde{U}^2}{\sqrt{\eta^2-1}}\sqrt{S_{k+1}}\cdot\begin{cases}\sin(\sum_{l=k}^{i-1}(-1)^{l-k}\alpha_l), & \text{if } \bar{N}\leq k\leq i\leq N-1;\\ \sin(\sum_{l=i}^{k-1}(-1)^{k-1-l}\alpha_l), & \text{if } \bar{N}\leq i< k\leq N-1.\end{cases} \tag{3.61}$$

All distance-1 interpolation inequalities $Q_{i,j}$ with $|i-j|=1$ are involved in the proofs and thus hold with equality. This fact can also be checked by direct substitutions of our triplets and using that the inner product between consecutive gradients is known in closed form; for $i \geq \bar{N}$ we have $\langle g_{\bar{N}}, g_{\bar{N}+1}\rangle = cU^2$, with $c = \frac{1+\eta\rho}{\eta+\rho}$ (as in (3.39)).

Case $i > j$. Let us consider a fixed $j \in \{0, 1, \dots, N-1\}$. For each $i \in \{j+1, \dots, N\}$, we define

$$\Delta Q_i \coloneqq Q_{i+1,j} - Q_{i,j} - Q_{i+1,i}.$$

The distance-1 interpolation inequalities hold exactly, hence $Q_{i+1,i} = 0$ and $Q_{j+1,j} = 0$. Therefore to prove that $Q_{i,j} \geq 0$ for all $i > j$ it is sufficient to show that $\Delta Q_i \geq 0$, as $Q_{i,j}$ is obtained as a sum of positive increments.
Substituting the expressions of $Q_{i+1,j}$, $Q_{i,j}$ and $Q_{i+1,i}$ we subsequently obtain:

$$\begin{aligned}
\Delta Q_i &= \frac{1}{L-\mu}\left\langle g_i - g_j - \mu(x_i - x_j)\, ,\, g_i - g_{i+1} - L(x_i - x_{i+1})\right\rangle \\
&= \frac{1}{L-\mu}\langle g_j - g_i - \gamma\mu \sum_{k=j}^{i-1} g_k\, ,\, g_{i+1} - (1-\gamma L) g_i\rangle \\
&= \frac{-\gamma}{\eta-\rho}\sum_{k=j}^{i-1}\left\langle g_{k+1} - \eta g_k\, ,\, g_{i+1} - \rho g_i\right\rangle,
\end{aligned}$$

where in the second step we use the gradient descent iteration. For any $k \leq \bar{N}$, from (3.58) we have that $\langle g_{k+1} - \eta g_k\, ,\, g_{i+1} - \rho g_i\rangle = 0$. Hence it is sufficient to focus on the range $j \geq \bar{N}+1$ (thus $i \geq \bar{N}+2$). Using (3.61) in the case $i \geq k$, the expression of ΔQ_i can be rewritten as

$$\Delta Q_i = \frac{\gamma(\eta-\rho)\tilde{U}^2}{\sqrt{\eta^2-1}}\sum_{k=j}^{i-1}\sqrt{S_{k+1}}\sin\Big(\sum_{l=k}^{i-1}(-1)^{l-k}\alpha_l\Big).$$

Because α_l decreases monotonically on $(0, \frac{\pi}{2})$ and the first term of the sum is always positive, then the sine of the sum is also positive, implying $\Delta Q_i \geq 0$.

Case $i < j$. Let us consider a fixed $j \in \{1, 2, \dots, N\}$. For each $i \in \{0, 1, \dots, j-1\}$, we define

$$\Delta Q_i \coloneqq Q_{i,j} - Q_{i+1,j} - Q_{i,i+1}.$$

The distance-1 interpolation inequalities hold exactly, hence $Q_{i+1,i} = 0$ and $Q_{j+1,j} = 0$. Therefore to prove $Q_{i,j} \geq 0$ for all $i < j$ it is sufficient to show that

$\Delta Q_i \geq 0$, since $Q_{i,j}$ is obtained as a sum of positive increments.
Similarly to the previous case, substituting the expressions of $Q_{i,j}$, $Q_{i+1,j}$ and $Q_{i,i+1}$, we subsequently obtain:

$$\Delta Q_i = \frac{1}{L-\mu} \left\langle g_{i+1} - g_j - \mu(x_{i+1} - x_j)\, ,\, g_{i+1} - g_i - L(x_{i+1} - x_i)\right\rangle$$

$$= \frac{-\gamma}{\eta - \rho} \sum_{k=i+1}^{j-1} \left\langle g_{k+1} - \eta g_k\, ,\, g_{i+1} - \rho g_i \right\rangle .$$

From (3.58) we have that $\Delta Q_i = 0$ if $j \leq \bar{N}+1$. Therefore, further on we assume $j \geq \bar{N}+2$. Moreover, we have the identity $g_{k+1} - \rho g_k = \eta^{\bar{N}-k}(g_{\bar{N}+1} - \rho g_{\bar{N}})$ for $k = 0, 1, \ldots, \bar{N}$, hence, using (3.58), ΔQ_i rewrites as

$$\Delta Q_i = \frac{-\gamma}{\eta-\rho}\eta^{\max\{0,\bar{N}-i\}} \sum_{k=\max\{i,\bar{N}\}+1}^{j-1} \left\langle g_{k+1} - \eta g_k\, ,\, g_{\max\{i,\bar{N}\}+1} - \rho g_{\max\{i,\bar{N}\}} \right\rangle .$$

From (3.61) in the case $i < k$ we get

$$\Delta Q_i = \frac{\gamma(\eta-\rho)\tilde{U}^2\eta^{\max\{0,\bar{N}-i\}}}{\sqrt{\eta^2-1}} \sum_{k=\max\{i,\bar{N}\}+1}^{j-1} \sqrt{S_{k+1}} \sin\left(\sum_{l=\max\{i,\bar{N}\}}^{k-1} (-1)^{k-1-l}\alpha_l \right) .$$

It is sufficient to check the positivity of the following expression for any $\bar{N} \leq i < j$

$$V := \sum_{k=i+1}^{j-1} \sqrt{S_{k+1}} \sin\Big(\sum_{l=i}^{k-1} (-1)^{l-(k-1)}\alpha_l \Big).$$

Let $\Sigma_k := \sum_{l=i}^{k-1}(-1)^{l-(k-1)}\alpha_l$ and $Z_k := \sqrt{S_{k+1}}\sin(\Sigma_k)$. For any integer $p \geq 0$, one can check the following properties:

1. $\Sigma_{(i+1)+2p} + \Sigma_{(i+1)+2p+1} = \alpha_{(i+1)+2p}$;
2. $\Sigma_{(i+1)+2p}$ is a decreasing sequence;
3. $\Sigma_{(i+1)+2p} \geq \alpha_{(i+1)+2p}$;
4. $\Sigma_{(i+1)+2p} \in (0, \frac{\pi}{2})$.

Then we have

$$Z_{(i+1)+2p} = \sqrt{S_{(i+1)+2p+1}}\sin(\Sigma_{(i+1)+2p}) > 0;$$

$$Z_{(i+1)+2p+1} = -\sqrt{S_{(i+1)+2p+2}}\sin(\Sigma_{(i+1)+2p} - \alpha_{(i+1)+2p}) < 0.$$

A sufficient condition for the positivity of V is $Z_{(i+1)+2p}+Z_{(i+1)+2p+1} \geq 0$ for any $p \geq 0$. This condition is equivalent to $\sqrt{\frac{S_{(i+1)+2p+1}}{S_{(i+1)+2p+2}}} > \frac{\sin(\Sigma_{(i+1)+2p}-\alpha_{(i+1)+2p})}{\sin(\Sigma_{(i+1)+2p})}$. Using the monotonic decrease of the cotangent function on $(0, \frac{\pi}{2})$ and the property $\Sigma_{(i+1)+2p} \leq \Sigma_{(i+1)} = \alpha_i$, we have that

$$\begin{aligned}\frac{\sin(\Sigma_{(i+1)+2p} - \alpha_{(i+1)+2p})}{\sin(\Sigma_{(i+1)+2p})} &= \sin(\alpha_{(i+1)+2p})(\cot(\alpha_{(i+1)+2p}) - \cot(\Sigma_{(i+1)+2p})) \\ &\leq \sin(\alpha_{(i+1)+2p})(\cot(\alpha_{(i+1)+2p}) - \cot(\alpha_i)).\end{aligned}$$

From (3.60) we have $\cot(\alpha_l) = \frac{-\eta\rho\sqrt{S_{l+1}}+\sqrt{S_{l+2}}}{\sqrt{\eta^2-1}}$,for any $l \geq \bar{N}$. Then

$$\begin{aligned}&\sin(\alpha_{(i+1)+2p})(\cot(\alpha_{(i+1)+2p}) - \cot\alpha_i) \\ &\qquad = \frac{-\eta\rho\left(\sqrt{\frac{S_{(i+1)+2p+1}}{S_{(i+1)+2p+2}}} - \sqrt{\frac{S_{i+1}}{S_{(i+1)+2p+2}}}\right) + \left(1 - \sqrt{\frac{S_{i+2}}{S_{(i+1)+2p+2}}}\right)}{\eta - \rho\sqrt{\frac{S_{(i+1)+2p+1}}{S_{(i+1)+2p+2}}}},\end{aligned}$$

where we amplified by $\frac{1}{\sqrt{S_{(i+1)+2p+2}}}$. Thus it is left to show that for each $p \geq 0$ it holds

$$\sqrt{\frac{S_{(i+1)+2p+1}}{S_{(i+1)+2p+2}}} \geq \frac{-\eta\rho\left(\sqrt{\frac{S_{(i+1)+2p+1}}{S_{(i+1)+2p+2}}} - \sqrt{\frac{S_{i+1}}{S_{(i+1)+2p+2}}}\right) + \left(1 - \sqrt{\frac{S_{i+2}}{S_{(i+1)+2p+2}}}\right)}{\eta - \rho\sqrt{\frac{S_{(i+1)+2p+1}}{S_{(i+1)+2p+2}}}},$$

which is equivalent to:

$$\begin{aligned}\eta(1+\rho) - \left(\rho + \frac{S_{(i+1)+2p+2}}{S_{(i+1)+2p+1}}\right)\sqrt{\frac{S_{(i+1)+2p+1}}{S_{(i+1)+2p+2}}} &\geq \\ &\eta\rho\sqrt{\frac{S_{i+1}}{S_{(i+1)+2p+1}}} - \sqrt{\frac{S_{i+2}}{S_{(i+1)+2p+1}}}.\end{aligned}$$

Using $\frac{S_{(i+1)+2p+2}}{S_{(i+1)+2p+1}} = \rho^2 + \frac{1}{S_{(i+1)+2p+1}}$, the inequality equivalently rewrites as:

$$(1+\rho)\left(\eta - \rho\sqrt{\frac{S_{(i+1)+2p+1}}{S_{(i+1)+2p+2}}}\right) + \frac{\frac{-1}{\sqrt{S_{(i+1)+2p+2}}} + \sqrt{S_{i+2}} - \eta\rho\sqrt{S_{i+1}}}{\sqrt{S_{(i+1)+2p+1}}} \geq 0.$$

The first term is positive because $\rho \in (-1, 0)$, whereas the second one due to $S_{i+2} \geq 1$ and $\frac{1}{\sqrt{S_{(i+1)+2p+2}}} < 1$. □

Chapter 4

Tight convergence rates for the difference-of-convex algorithm and proximal gradient descent

Based on

[101] Teodor Rotaru, Panagiotis Patrinos, and François Glineur. Tight analysis of difference-of-convex algorithm (DCA) improves convergence rates for proximal gradient descent. In *The 28th International Conference on Artificial Intelligence and Statistics*, volume 258, pages 4114–4122, 2025

[102] Teodor Rotaru, Panagiotis Patrinos, and François Glineur. Tight convergence rates in gradient mapping for the difference-of-convex algorithm. In Hoai An Le Thi, Tao Pham Dinh, and Hoai Minh Le, editors, *Modelling, Computation and Optimization in Information Systems and Management Sciences*, pages 61–73, Cham, 2026. Springer Nature Switzerland

The results in Section 4.4 belong to the unpublished preprint [103]. The results in Section 4.6 were only partially published before in both papers [101, 102]. The results in Section 4.2.3 were not published before.

4.1 Introduction

In Chapter 3 we considered the unconstrained setting for smooth functions, that can be nonconvex or (strongly) convex, and provided a comprehensive analysis of the gradient descent method with constant stepsizes. The next natural step is to include constraints, which can be handled by the projected gradient descent method or, more generally, by the proximal gradient descent (PGD) method. The analysis of proximal gradient descent introduces two additional curvature bounds in comparison to gradient descent, corresponding to the function that incorporates the constraints.

However, this added difficulty can be simplified by exploiting a lesser-known connection from the literature: one iteration of the PGD method is exactly equivalent to one iteration of the difference-of-convex algorithm (DCA), when considering a suitable splitting of the objective function. DCA is a parameter-free method, meaning that it does not require a stepsize selection, and is therefore more suitable for performing a comprehensive worst-case analysis.

In this chapter, we focus on deriving a tight characterization of DCA under two different performance measures. To better understand the behaviour of PGD for different parameter choices, we extend the standard DCA setting to include possibly weakly convex (or hypoconvex) functions as the subtracted components. Keeping in mind the equivalence with PGD, and although DCA is primarily designed for nonconvex problems, we also allow strongly convex or concave objectives in our analysis.

Consider the difference-of-convex (DC) formulation

$$\underset{x\in\mathbb{R}^d}{\text{minimize}}\, F(x) \coloneqq f_1(x) - f_2(x), \tag{DC}$$

where $f_1, f_2 \colon \mathbb{R}^d \to \mathbb{R}$ are proper, lower semicontinuous convex functions, and F is lower bounded.

A standard method to solve (DC) is the difference-of-convex algorithm (DCA), a versatile method with no parameter that can find a *critical* point of F, defined as a point $x^* \in \mathbb{R}^d$ for which there exist subgradients $g_1^* \in \partial f_1(x^*)$ and $g_2^* \in \partial f_2(x^*)$ such that $g_1^* = g_2^*$. Stationary points of F are always critical, but the converse is not true. Extensive analyses of DCA are provided by Dinh and Thi [39], Tao and An [109], Horst and Thoai [62], Le Thi and Pham Dinh [74] and [38]. DCA is also referred to as the convex-concave procedure (CCCP), as seen in the work of Yuille and Rangarajan [135], Lanckriet and Sriperumbudur [71], Lipp and Boyd [79]. Interestingly, convergence analysis of many methods can be reduced to the one of DCA; for example, the Frank-Wolfe algorithm [49], as proved by Yurtsever and Sra [136], or the proximal gradient descent

(PGD), as shown by Le Thi and Pham Dinh [74, §3.3.4]. Conversely, Faust et al. [48] show that DCA is an instance of the Bregman proximal point algorithm. Linear convergence can be shown under the Polyak-Łojasiewicz (PL) inequality [72, 133].

An extensive list of DCA applications is provided by Le Thi and Pham Dinh [74]. Notable examples include efficient formulations for clustering problems (Hoai An et al. [61]), dictionary learning (Vo et al. [127]), robust support vector regression (Wang et al. [129]), multi-class support vector machines (MSVM) (Le Thi and Nguyen [73]), sparse logistic regression (Yang and Qian [132]), compressed sensing (Yin et al. [134]), adversarial attack for adversarial robustness and approximate optimization of complex functions (Awasthi et al. [9]) or Shallow Multilayer Perceptron (MLP) Neural Networks (Askarizadeh et al.[7]). Sun et al. [107] introduce the Negative ResNets, where f_2 is weakly convex.

Abbaszadehpeivasti et al. [3] provide exact convergence rates of DCA when both functions are convex. Their approach is based on performance estimation, including rigorous proofs for the rates, their exactness being supported by strong numerical evidence and, in some cases, by the identification of instances matching those rates exactly.

In this chapter, we study a generalization of the standard (DC) setting, considering the case where f_2 can be weakly convex (or hypoconvex). Some previous works also introduce weak convexity in either f_1 (Sun and Sun [106]) or f_2 (Syrtseva et al. [108]).

A key motivation for examining the case with f_2 weakly convex is that it mirrors the behaviour of applying PGD with stepsizes larger than the inverse Lipschitz constant (see Section 4.6). Additionally, our generalized DCA framework provides a useful tool for analyzing exact rates for PGD, with the benefit of handling one fewer parameter – DCA involves four curvature parameters compared to PGD's four curvature parameters plus the stepsize. Therefore, due to the equivalence of the iterations, it is more convenient to use a DCA-like analysis.

4.1.1 Contributions

Our results follow the same line as Abbaszadehpeivasti et al. [3], also relying on performance estimation.

When one of the functions is smooth, we provide an exact characterization of a single (DCA) iterate, measuring either the *(best) residual gradient* in

Theorem 4.2.1 or the *(best) (composite) gradient mapping* (also referred to as *iterate progress*) in Theorem 4.3.1. These metrics are rigorously defined below. Each of the two theorems describes a total of six distinct regimes, partitioning the parameter space according to the smoothness and strong convexity properties of the two component functions. Among these regimes, only two were previously known and proved by Abbaszadehpeivasti et al. [3], corresponding to the standard (DCA) setting where both f_1 and f_2 are convex, and where the objective F is assumed to be neither convex nor concave. Our analysis refines those earlier proofs by providing explicit descent lemmas.

We show that DCA converges sublinearly to critical points, with an $\mathcal{O}(1/N)$ rate after N iterations, again with six distinct regimes. Based on strong numerical evidence, we conjecture that three of those rates are exact for any number of iterations. For the remaining three, building on insights from gradient descent developed in Chapter 3, we provide their exact behaviour across all curvature configurations; some regimes are rigorously proven, while others remain as conjectures.

In the case where both functions are nonsmooth, we derive an $\mathcal{O}(1/N)$ sublinear rate in Theorem 4.4.2, strengthening and generalizing the existing results of [3] by employing an appropriately adapted optimality criterion.

Leveraging the worst-case performance analysis, we show that a decomposition of the objective F allowing weak convexity of f_2 can yield better rates compared to the standard DCA. Moreover, when both functions are smooth, a well-chosen DC splitting may even surpass the classical gradient descent method.

As a direct consequence of our detailed analysis, in Section 4.6 the convergence rates obtained in specific regimes are readily transferred to the PGD setting for any constant stepsize, including the ones exceeding the inverse Lipschitz constant.

Finally, we propose a general adaptive curvature shifting procedure for DCA, leveraging PGD stepsize schedules that align with established gradient descent strategies.

Summary of main results

Our contributions relative to the state of the art are outlined in Table 4.1.1, while Table 4.1.2 provides a comprehensive roadmap of our theoretical contributions, organized by performance metric and result type.

Proved contributions:

Table 4.1.1: State of the art vs contributions for difference-of-convex algorithm analysis. Legend: **Bold** = complete/proven results in this chapter. Abbreviations: w.c./c./s.c. = weakly convex / convex / strongly convex.

Criterion	State of the art	This chapter
Function class	Standard DC: f_1, f_2 convex [3]; limited extensions to w.c. f_2 [108]	**Complete**: f_1 convex or s.c., f_2 w.c./c./s.c.; $f_1 \in \mathcal{F}_{\mu_1,L_1}$, $f_2 \in \mathcal{F}_{\mu_2,L_2}$
Performance metrics	Residual gradient for f_1, f_2 convex [3]; Bregman divergence [3, 48]	**Three unified metrics**: residual gradient, gradient mapping, Bregman divergence across all regimes
Proved upper bounds	Standard DC ($\mu_1, \mu_2 \geq 0$): two regimes for residual gradient and two for gradient mapping [3]; Bregman divergence [3, 48]	**Six regimes** for residual gradient and six for gradient mapping; extended Bregman divergence bounds for w.c. f_2
Proved tightness	Two regimes tight for standard DC [3]	**Two residual gradient regimes proved tight** (any N); four additional conjectured as tight for one iteration; gradient mapping regimes conjectured tight for one iteration
Proof technique	PEP-based certificates [3]; no explicit descent lemmas	**Explicit descent lemmas** for all regimes in both metrics; algebraic proofs
Extension beyond one iteration	Sublinear rates for standard DC [3]	**Tight conjectures** for arbitrary N; proved linear rates for certain subdomains
PGD equivalence	Known connection [74]; limited rate transfers	**Complete rate transfer**: all DCA regimes translated to PGD with any constant $\gamma \in (0, \frac{2}{L_\varphi})$
Curvature shifting	Standard practice: shift to convexity [74]	**Optimal shifting**: weak convexity of f_2 can improve rates; numerical validation

i) **Extension to weakly convex** f_2: We introduce new regimes (p_3, p_5, r_3, r_5) allowing $\mu_2 < 0$, thereby accommodating weakly convex functions in the extended DC splitting.

Table 4.1.2: (DCA) Overview of main results for DCA and PGD analysis

Result type	Performance metric	Key results	Status	Section
One-step descent	Residual gradient	Theorem 4.2.1 (6 regimes p_1–p_6)	✓	4.2.1
One-step descent	Gradient mapping	Theorem 4.3.1 (6 regimes r_1–r_6)	✓	4.3.1
One-step descent	Bregman divergence	Theorem 4.4.1 (both nonsmooth)	✓	4.4
Sublinear rates	Residual gradient	Corollary 4.2.1 ($\mathcal{O}(1/N)$)	✓	4.2.1
Sublinear rates	Gradient mapping	Theorem 4.3.1 ($\mathcal{O}(1/N)$)	✓	4.3.1
Tightness	After 1 iteration	Proposition 4.2.1 (p_1, p_2); Conjectures 4.2.1 and 4.2.2	Mixed	4.2.2
Linear rates	Residual gradient	Theorems 4.2.2 and 4.2.3	✓	4.2.3
Linear rates	Gradient mapping	Theorems 4.3.2 and 4.3.3	✓	4.3.2
Tight rates (general)	Both metrics	Conjectures 4.2.3 to 4.2.5 and 4.3.1	⋆	4.2.3, 4.3.2
PGD equivalence	Both metrics	Propositions 4.6.1 and 4.6.2; Corollaries 4.6.2 to 4.6.5	Mixed	4.6

Legend: ✓ = Proved; ⋆ = Conjectured with strong PEP numerical evidence; Mixed = Some parts proved, others conjectured.

ii) **Explicit descent lemmas**: For all six regimes in both residual gradient and gradient mapping metrics, we provide complete algebraic proofs via weighted combinations of interpolation inequalities, improving upon previous PEP-based implicit characterizations.

iii) **Linear convergence in specific subdomains**: For strongly convex objective F, using distance-1 interpolation inequalities, we establish exact linear rates in Theorems 4.2.2, 4.2.3, 4.3.2 and 4.3.3.

iv) **PGD-DCA equivalence and rate transfer**: We fully characterize the iteration-wise equivalence (Propositions 4.6.1 and 4.6.2), enabling complete convergence rate transfer for *any* constant stepsize $\gamma \in (0, \frac{2}{L_\varphi - \mu_h})$, including previously uncharacterized long stepsizes $\gamma > \frac{1}{L_\varphi}$.

v) **Curvature shifting optimization framework**: We develop a systematic approach for optimal selection of the curvature parameter λ in the splitting $F = (f_1 - \lambda\frac{\|\cdot\|^2}{2}) - (f_2 - \lambda\frac{\|\cdot\|^2}{2})$, demonstrating via sparse PCA experiments that allowing weak convexity in f_2 can significantly improve practical convergence rates (Section 4.7).

Conjectured contributions (strongly supported by PEP):

i) **Complete tightness characterization**: Conjectures 4.2.1 and 4.2.2 establish the tightness of all six regimes, validated through extensive PEP numerical experiments spanning broad parameter ranges (see GitHub repository for reproducibility).

ii) **Exact rates for arbitrary** N: Conjectures 4.2.3 to 4.2.5 and 4.3.1 provide closed-form expressions for tight convergence rates after any number of iterations, directly inspired by and extending the exact gradient descent analysis in Chapter 3.

We provide a GitHub repository[1] to support the numerical conjectures and to reproduce all simulations for the residual gradient results.

4.1.2 Theoretical background

Assumption 4.1.1 (Objective and parameters)**.** *The objective function F in (DC) is lower bounded and can be written as $F = f_1 - f_2$, where $f_1 \in \mathcal{F}_{\mu_1,L_1}$ and $f_2 \in \mathcal{F}_{\mu_2,L_2}$, with parameters $\mu_1 \in [0,\infty)$, $L_1 \in (0,\infty]$, $\mu_2 \in (-\infty,\infty)$ and $L_2 \in (\mu_2,\infty]$, such that $\mu_1 < L_1$ and $\mu_2 < L_2$.*

Assumption 4.1.1 runs throughout the rest of this chapter and implies $F \in \mathcal{F}_{\mu_1-L_2\,,\,L_1-\mu_2}$. The function f_2 may be weakly convex or even concave, the latter case corresponding to analyzing the PGD iteration on strongly convex functions with long stepsizes (see Section 4.6). We also denote $F_{lo} \coloneqq \inf_x F$.

DCA iteration. With $g_2 \in \partial f_2(x)$, one DCA iteration selects

$$x^+ \in \operatorname*{arg\,min}_{w\in\mathbb{R}^d}\{f_1(w) - \langle g_2, w\rangle\}. \tag{DCA}$$

With an abuse of notation, a more compact definition of the (DCA) iteration is $x^+ \in \partial f_1^*(\partial f_2(x))$. We state the method in Algorithm 2.

[1] Repository: https://github.com/teo2605/DCA_AISTATS25

Algorithm 2: Difference-of-convex algorithm (DCA)

Data: $f_1 \in \mathcal{F}_{\mu_1,L_1}$, $f_2 \in \mathcal{F}_{\mu_2,L_2}$, with $\mu_1 \geq 0$, $\mu_1 + \mu_2 > 0$ or $\mu_1 = \mu_2 = 0$; $N \geq 1$ iterations starting from $x^0 \in \operatorname{dom} \partial f_1$

```
for k = 0, ..., N do
    Select g_2^k ∈ ∂f_2(x^k) ;
    Select x^{k+1} ∈ argmin_{w∈R^d} {f_1(w) − ⟨g_2^k, w⟩}
```

The optimality condition in the definition of x^+ implies the existence of $g_1^+ \in \partial f_1(x^+)$ such that $g_1^+ = g_2$, where $g_2 \in \partial f_2(x)$. This is the only characterization of x^+ used in our derivations. Note that the sequence of iterates produced by (DCA) is not unique.

Assumption 4.1.2 (Subdifferential domain-range compatibility)**.** *The subdifferentials of f_1 and f_2 satisfy the following conditions: $\varnothing \neq \operatorname{dom} \partial f_1 \subseteq \operatorname{dom} \partial f_2$ and $\operatorname{range} \partial f_2 \subseteq \operatorname{range} \partial f_1$.*

Proposition 4.1.1 (Well-definiteness of DCA iterations)**.** *Under Assumption 4.1.2, the (DCA) iterations are well-defined, meaning there exists a sequence $\{x^k\}$, starting from $x^0 \in \operatorname{dom} \partial f_1$, generated by $x^{k+1} \in \partial f_1^*(\partial f_2(x^k))$.*

Tao and An [109] note that DCA is typically well-defined, as for any l.s.c. function f, it holds $\operatorname{ri}(\operatorname{dom} f) \subseteq \operatorname{dom} \partial f \subseteq \operatorname{dom} f$, where $\operatorname{ri}(\operatorname{dom} f)$ is the relative interior of $\operatorname{dom} f$. The potential weak convexity of f_2 represents only a curvature adjustment in the subdifferential definition.

A critical point x^* satisfies $\partial f_2(x^*) \cap \partial f_1(x^*) \neq \emptyset$. When both functions f_1 and f_2 are smooth, any critical point x^* is clearly stationary as we have $\nabla F(x^*) = \nabla f_1(x^*) - \nabla f_2(x^*) = 0$. If only f_2 is smooth, we have $\partial F(x^*) = \partial f_1(x^*) - \nabla f_2(x^*)$ (Rockafellar and Wets [98, Exercise 10.10]) and criticality also implies stationarity, since $0 \in \partial F(x^*)$. However, if only f_1 is smooth, we can only guarantee the inclusion $\partial(-f_2)(x^*) \subseteq -\partial f_2(x^*)$ (Rockafellar and Wets [98, Corollary 9.21]), implying only $\partial F(x^*) \subseteq \nabla f_1(x^*) - \partial f_2(x^*)$, and critical points may not be stationary.

If the entire sequence of iterates $\{x^k\}$ is bounded, then each accumulation point of it is a critical point of F [39, Theorem 3(iv)].

Proposition 4.1.2 (DCA sufficient decrease)**.** *Let $f_1 \in \mathcal{F}_{\mu_1,L_1}$ and $f_2 \in \mathcal{F}_{\mu_2,L_2}$. If $\mu_1 + \mu_2 \geq 0$, then the objective function $F = f_1 - f_2$ decreases after each iteration of (DCA). Moreover, if $\mu_1 + \mu_2 > 0$ that objective decrease is strict, unless $x^+ = x$.*

Proof. Using Lemma 2.2.1 on function f_1 for the pair $(x, y) = (x, x^+)$ and on function f_2 with $(x, y) = (x^+, x)$, and summing the inequalities, we obtain:

$$\frac{\mu_1 + \mu_2}{2}\|x^+ - x\|^2 \leq f_1(x) - f_1(x^+) - \langle g_1^+\,,\, x - x^+\rangle +$$

$$f_2(x^+) - f_2(x) - \langle g_2\,,\, x^+ - x\rangle \leq \frac{L_1 + L_2}{2}\|x^+ - x\|^2,$$

where $g_1^+ \in \partial f_1(x^+)$ and $g_2 \in \partial f_2(x)$. By using $F(x) = f_1(x) - f_2(x)$ and $g_1^+ = g_2$ we get:

$$\frac{\mu_1 + \mu_2}{2}\|x - x^+\|^2 \leq F(x) - F(x^+) \leq \frac{L_1 + L_2}{2}\|x - x^+\|^2,$$

enough to prove both parts of the proposition. □

The proof of Proposition 4.1.2 is inspired by [39, Theorem 3, Proposition 2].

Remark 4.1.1 (Exact oracle assumption)**.** *Throughout this chapter, we assume available the exact oracles of ∂f_1 and ∂f_1^*, thus we only focus on the progress of the iterations. This assumption is significant when comparing the improvements obtained via curvature shifting.*

Performance measures. When at least one of the functions is smooth, we track one of the following two quantities:

i) *(Best) residual gradient norm*

$$\min_{0\leq k\leq N}\{\|g_1^k - g_2^k\|^2\},$$

where $g_i \in \partial f_i(x^k)$, $i = \{1, 2\}$, are any subgradients of f_i evaluated at x^k;

ii) *(Best) (composite) gradient mapping* [88, 89] (also known as the *prox-gradient mapping* [36] or (best) iterate progress)

$$\min_{0\leq k\leq N}\{\|x^k - x^{k+1}\|^2\}.$$

In each case we get six regimes, denoted by p_i, with $i = \{1, \ldots, 6\}$.
When both functions are nonsmooth, we employ a different measure based on Bregman divergence (for details, see Section 4.4).

Notation. Specifically to this chapter, superscripts indicate the iteration index (e.g., x^k represents the k-th iterate). We define $E_k\colon \mathbb{R}_{>0} \to \mathbb{R}_{>0}$,

$$E_k(z) := \sum_{j=1}^{2k} z^{-j} = \begin{cases} \frac{-1+z^{-2k}}{1-z}, & \text{if } z \neq 1; \\ 2k, & \text{if } z = 1. \end{cases} \tag{E_k}$$

Note that $E_0(z) = 0$.

Chapter organization. The performance bounds measuring the residual gradient norm, gradient mapping and Bregman divergence are presented in Section 4.2, Section 4.3 and Section 4.4, respectively. All corresponding proofs are provided in Section 4.5. The equivalence of PGD with DCA, along with the resulting convergence rates, is detailed in Section 4.6. Finally, the curvature shifting technique is discussed in Section 4.7, together with numerical experiments showing its benefits.

4.2 Performance bounds in residual gradient

This section focuses on deriving and discussing performance bounds for the residual gradient criterion, assuming that at least one of the two functions is smooth. In Section 4.2.1, we present a single-iteration analysis leading to six regimes yielding sublinear rates. Their tightness is discussed in Section 4.2.2, while tight bounds for an arbitrary horizon, some of which remain conjectural, are provided in Section 4.2.3.

4.2.1 Sublinear convergence rates in residual gradient

Theorem 4.2.1 (DCA one-step descent for the residual gradient)**.** *Let $f_1 \in \mathcal{F}_{\mu_1,L_1}$ and $f_2 \in \mathcal{F}_{\mu_2,L_2}$ satisfy Assumptions 4.1.1 and 4.1.2, with at least f_1 or f_2 smooth, and assume $\mu_1 + \mu_2 > 0$ or $\mu_1 = \mu_2 = 0$. Then after one step of* (DCA) *we have*

$$F(x) - F(x^+) \geq \sigma_i \tfrac{1}{2}\|g_1 - g_2\|^2 + \sigma_i^+ \tfrac{1}{2}\|g_1^+ - g_2^+\|^2 \tag{4.1}$$

with $g_1 \in \partial f_1(x)$, $g_1^+ \in \partial f_1(x^+)$, $g_2 \in \partial f_2(x)$, $g_2^+ \in \partial f_2(x^+)$, and the expressions for $\sigma_i, \sigma_i^+ \geq 0$ correspond to one of the six regimes (indexed by $i = 1, \ldots, 6$) described in Table 4.2.1 according to the values of parameters L_1, L_2, μ_1, μ_2.

The six regimes appearing in Table 4.2.1 are illustrated in Figure 4.2.1; we refer to each p_i as one of the six regimes together with its corresponding expression.

Table 4.2.1: Exact decrease in residual gradient norm after one iteration: $F(x) - F(x^+) \geq \sigma_i \frac{1}{2}\|g_1 - g_2\|^2 + \sigma_i^+ \frac{1}{2}\|g_1^+ - g_2^+\|^2$ (see Theorem 4.2.1), with $\sigma_i, \sigma_i^+ \geq 0$ and $p_i = \sigma_i + \sigma_i^+$, $i = 1, \ldots, 6$. The domains satisfy the condition $\mu_1 + \mu_2 > 0$ or $\mu_1 = \mu_2 = 0$, with at least one between f_1 and f_2 smooth. Notation: $B \coloneqq \mu_1^{-1} + \mu_2^{-1} + L_2^{-1}$ and $E \coloneqq \frac{L_2+\mu_2}{L_1 L_2}\frac{L_2-L_1}{-\mu_2} + \mu_1^{-1} - L_1^{-1}$. (s.c.: strongly convex; n.c.: nonconvex; ccv.: concave; n.ccv.: nonconcave)

Regime	σ_i	σ_i^+	**Domain**			**Description**
p_1	$L_2^{-1}\frac{L_2-\mu_1}{L_1-\mu_1}$	$L_2^{-1}\left(1+\frac{L_2^{-1}-L_1^{-1}}{\mu_1^{-1}-L_1^{-1}}\right)$	$L_1 \geq L_2 \geq \mu_1 \geq 0$ $L_1 > \mu_2$	$\mu_2 \geq 0$		f_1, f_2 convex F n.c.-n.ccv.
				$\mu_2 < 0,\ E \leq 0$		f_1 s.c., f_2 n.c. F n.c.-n.ccv.
p_2	$L_1^{-1}\left(1+\frac{L_1^{-1}-L_2^{-1}}{\mu_2^{-1}-L_2^{-1}}\right)$	$L_1^{-1}\frac{L_1-\mu_2}{L_2-\mu_2}$	$L_2 \geq L_1 \geq \mu_2 \geq 0$ $L_2 > \mu_1$	$\mu_1 \geq 0$		f_1, f_2 convex F n.c.-n.ccv.
p_3	$\frac{L_1^{-1}(\mu_1^{-1}+\mu_2^{-1}+L_2^{-1})}{\mu_1^{-1}+\mu_2^{-1}+L_2^{-1}-L_1^{-1}}$	$\frac{1}{L_2+\mu_2}$	$\mu_2 < 0,\ \mu_1 > 0$ $L_2 > \mu_1;\ L_1 > \mu_2$	$B \leq 0$	$L_1 \geq L_2,\ E \geq 0$	f_1 s.c., f_2 n.c. F n.c.-n.ccv.
					$L_2 > L_1$	
p_4	0	$\frac{L_2+\mu_1}{L_2^2}$	$L_1 > \mu_1 \geq L_2 > 0$ $L_1 > \mu_2$	$\mu_2 \geq 0$		f_1 s.c., f_2 convex F convex
				$\mu_2 < 0,\ B \leq 0$		f_1 s.c., f_2 n.c. F convex
p_5	0	$\frac{\mu_1+\mu_2}{\mu_2^2}$	$\mu_2 < 0,\ \mu_1 > 0$ $L_1 > \mu_2$	$B > 0$	$L_2 > \mu_1 > 0$	f_1 s.c., f_2 n.c. F n.c.-n.ccv.
					$0 < L_2 \leq \mu_1$	f_1 s.c., f_2 n.c. F convex
				$B \leq 0$	$L_2 \leq 0$	f_1 s.c., f_2 ccv. F s.c.
p_6	$\frac{L_1+\mu_2}{L_1^2}$	0	$L_2 > \mu_2 \geq L_1 > \mu_1$	$\mu_1 \geq 0$		f_1 convex, f_2 s.c. F ccv.

Notably, there is a striking symmetry between regimes p_1 and p_2, as well as between p_4 and p_6. Specifically, the formulas for p_2 and p_6 in Table 4.2.1 can be derived from those of p_1 and p_4 by swapping $L_1 \iff L_2$, $\mu_1 \iff \mu_2$, and $\sigma_i \iff \sigma_i^+$. The proof of Theorem 4.2.1 is deferred to Section 4.5.1.

Figure 4.2.2 shows a decision tree outlining how to identify the corresponding regime from Theorem 4.2.1 based on the given curvature parameters.

Corollary 4.2.1 (DCA sublinear rates for the residual gradient)**.** *Let $f_1 \in \mathcal{F}_{\mu_1,L_1}$ and $f_2 \in \mathcal{F}_{\mu_2,L_2}$ satisfying Assumptions 4.1.1 and 4.1.2, assume at least f_1 or f_2 is smooth, and assume $\mu_1 + \mu_2 > 0$ or $\mu_1 = \mu_2 = 0$. Then after N iterations of* (DCA) *starting from x^0 we have*

$$\tfrac{1}{2} \min_{0 \leq k \leq N} \{\|g_1^k - g_2^k\|^2\} \leq \frac{F(x^0) - F(x^N)}{p_i(L_1, L_2, \mu_1, \mu_2)N}, \tag{4.2}$$

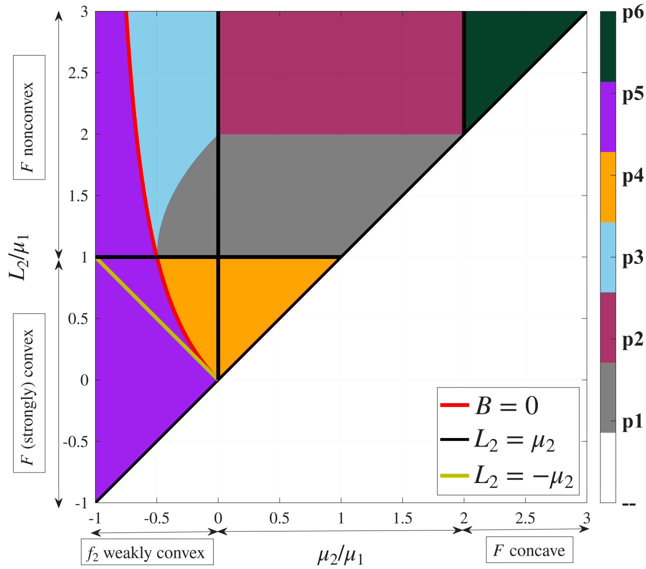

Figure 4.2.1: All regimes' domains when measuring residual gradient after one (DCA) iteration (Theorem 4.2.1), with $\mu_1 > 0$ and $^{L_1}/_{\mu_1} = 2$. Regimes are bounded by the conditions $L_2 > \mu_2$ and $\mu_2 > -\mu_1$. Regimes p_1, p_2 and p_3 lie within the area delimited by the *threshold* $B = 0$ (4.3) and the conditions $L_2 > \mu_1$ and $L_1 = 2\mu_1 > \mu_2$ (namely F is nonconvex-nonconcave), and are conjectured to be tight after N iterations of (DCA). For the rest, tight expressions are given in Section 4.2.3, proved in the areas of p_4 and p_5 (for $\mu_2 \leq -L_2$), and conjectured otherwise.

where $g_1^k \in \partial f_1(x^k)$ and $g_2^k \in \partial f_2(x^k)$ for all $k = 0, \ldots, N$ and $p_i = \sigma_i + \sigma_i^+$ is given in Table 4.2.1. Additionally, if F is nonconcave (i.e., $L_1 > \mu_2$):

$$\tfrac{1}{2} \min_{0 \leq k \leq N} \{\|g_1^k - g_2^k\|^2\} \leq \frac{F(x^0) - F_{lo}}{p_i(L_1, L_2, \mu_1, \mu_2)N + \frac{1}{L_1 - \mu_2}}.$$

Regimes p_1 and p_2 correspond in part to the standard setting of (DCA), where both functions are convex ($\mu_1 \geq 0$, $\mu_2 \geq 0$). Whether p_1 or p_2 holds depends on which is larger among L_1 and L_2. These regimes require the objective $F \in \mathcal{F}_{\mu_1 - L_2,\ L_1 - \mu_2}$ to be nonconvex ($L_2 > \mu_1$) and nonconcave ($L_1 > \mu_2$), and were first established by Abbaszadehpeivasti et al. [3], using performance estimation. All other described regimes are novel.

Remark 4.2.1 (Recovering known results for the standard convex case)**.** *In the specific convex scenario $\mu_1 = \mu_2 = 0$, both regimes p_1 and p_2 hold, as outlined*

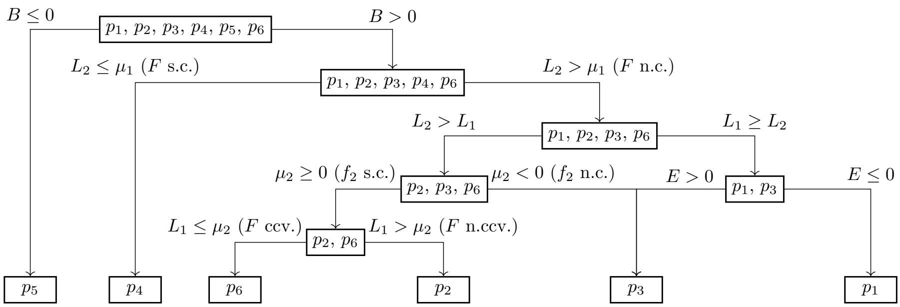

Figure 4.2.2: Decision tree on selecting regimes after one iteration defined in Theorem 4.2.1. Abbreviations – s.c.: strongly convex; n.c.: nonconvex; ccv.: concave; n.ccv.: nonconcave. Recall that $F = f_1 - f_2$, where we assume $f_1 \in \mathcal{F}_{\mu_1, L_1}$ and $f_2 \in \mathcal{F}_{\mu_2, L_2}$, with $\mu_1 + \mu_2 > 0$ or $\mu_1 = \mu_2 = 0$, and use the notation: $B \coloneqq \mu_1^{-1} + \mu_2^{-1} + L_2^{-1}$ and $E \coloneqq \frac{L_2+\mu_2}{L_1 L_2} \frac{L_2 - L_1}{-\mu_2} + \mu_1^{-1} - L_1^{-1}$.

by Abbaszadehpeivasti et al. [3, Corollary 3.1] and the one-step decrease is given by: $F(x) - F(x^+) \geq \frac{1}{2L_1}\|g_1 - g_2\|^2 + \frac{1}{2L_2}\|g_1^+ - g_2^+\|^2$. The same result is obtained using the Bregman proximal point algorithm perspective by Faust et al. [48, §4.2].

If f_1 is strongly convex, Theorem 4.2.1 actually extends regime p_1 beyond the difference-of-convex case, i.e., to situations where f_2 is weakly convex, such that $\mu_1 > 0 > -\mu_2$. This is valid up to a certain threshold determined by the sign of $E \coloneqq \frac{L_2+\mu_2}{L_1 L_2} \frac{L_2 - L_1}{-\mu_2} + \mu_1^{-1} - L_1^{-1}$. For p_1, the condition $E < 0$ holds, while for $E \geq 0$ regime p_3 emerges. Moreover, for $L_2 \geq L_1$ it always holds $E \geq 0$. Additionally, the boundary of regime p_3 is constrained by the threshold $B \leq 0$, where

$$B \coloneqq \mu_1^{-1} + \mu_2^{-1} + L_2^{-1}. \tag{4.3}$$

This condition corresponds to $\mu_1 > 0$ and $L_2 \leq \frac{-\mu_1 \mu_2}{\mu_1 + \mu_2}$. Regime p_4, emerging for $L_2 B > 0$, includes two cases: (i) when F is nonconvex-nonconcave; and (ii) when F is strongly convex (even containing f_2 concave with $L_2 \leq 0$). The threshold condition $B = 0$ (depicted by the red curve from Figure 4.2.1) distinguishes regime p_5 from p_3 and p_4. The latter are separated by the condition $L_2 = \mu_1$, delineating the cases F nonconvex (for p_3) and F (strongly) convex (for p_4), respectively. For completeness of analysis, we also include regime p_6, arising for a (strongly) concave objective (and unbounded from below), with $\mu_2 \geq L_1$.

For the particular setup $L_1 = 2$ and $\mu_1 = 1$, in Figure 4.2.1 we show the regions of regimes depending on curvatures L_2 and μ_2.

Table 4.2.2: Exact decrease in residual gradient after one iteration, with f_1 or f_2 nonsmooth: $F(x) - F(x^+) \geq \sigma_i \frac{1}{2}\|g_1 - g_2\|^2 + \sigma_i^+ \frac{1}{2}\|g_1^+ - g_2^+\|^2$, where $\sigma_i, \sigma_i^+ \geq 0$, $p_i = \sigma_i + \sigma_i^+$.

Regime	σ_i	σ_i^+	**Domain**
$p_{1,4}$	0	$\frac{L_2+\mu_1}{L_2^2}$	$L_1 = \infty > L_2 \geq \mu_1 \geq 0$ $\mu_2\left(\mu_1^{-1} + \mu_2^{-1} + L_2^{-1}\right) \geq 0$
$p_{2,6}$	$\frac{L_1+\mu_2}{L_1^2}$	0	$L_2 = \infty > L_1 \geq \mu_2 \geq 0$ $\mu_1\left(\mu_1^{-1} + \mu_2^{-1} + L_1^{-1}\right) \geq 0$
p_3	$\dfrac{\frac{1}{L_1}\left(\frac{1}{\mu_1}+\frac{1}{\mu_2}\right)}{\frac{1}{\mu_1}+\frac{1}{\mu_2}-\frac{1}{L_1}}$	0	$L_2 = \infty;\ \mu_1 > -\mu_2 > 0$
p_5	0	$\frac{\mu_1+\mu_2}{\mu_2^2}$	$L_1 = \infty;\ \mu_1 > -\mu_2 > 0$ $0 < \mu_1^{-1} + \mu_2^{-1} + L_2^{-1}$

One nonsmooth term

All the above results hold when at least one of the functions f_1 and f_2 is smooth. When exactly one of them is smooth, i.e., when the other is nonsmooth, some expressions in Table 4.2.1 become simpler, and we give an explicit description below. In the standard use of DCA, the conjugate step is applied to f_1 nonsmooth.

Corollary 4.2.2 (DCA sublinear rates for the residual gradient with one nonsmooth term)**.** *Let $f_1 \in \mathcal{F}_{\mu_1,L_1}$ and $f_2 \in \mathcal{F}_{\mu_2,L_2}$, where exactly one function f_1 or f_2 is smooth, and assume $\mu_1 + \mu_2 > 0$ or $\mu_1 = \mu_2 = 0$. Consider N iterations of* (DCA) *starting from x^0. Then:*

$$\tfrac{1}{2} \min_{0 \leq k \leq N} \{\|g_1^k - g_2^k\|^2\} \leq \frac{F(x^0) - F(x^N)}{p_i(L_1, L_2, \mu_1, \mu_2)N},$$

where $g_1^k \in \partial f_1(x^k)$ and $g_2^k \in \partial f_2(x^k)$ for all $k = 0, \ldots, N$ and p_i is provided in Table 4.2.2. Additionally, if F is nonconcave (i.e., $L_1 > \mu_2$):

$$\tfrac{1}{2} \min_{0 \leq k \leq N} \{\|g_1^k - g_2^k\|^2\} \leq \frac{F(x^0) - F_{lo}}{p_i(L_1, L_2, \mu_1, \mu_2)N + \frac{1}{L_1 - \mu_2}}.$$

Corollary 4.2.2 is derived by setting $L_1 = \infty$ or $L_2 = \infty$ in Corollary 4.2.1 and in the corresponding entries from Table 4.2.1. It shows identical rates as Abbaszadehpeivasti et al. [3, Corollary 3.1] for $\mu_1, \mu_2 \geq 0$, while extending them to scenarios involving weakly convex f_2, up to the threshold $B = 0$, beyond which regime p_5 emerges. Notably, we observe in Table 4.2.2 that regimes p_4 and p_6 condense to regimes p_1 and p_2, respectively. When $L_1 = \infty$, only regimes

p_1 and p_5 hold, separated by the threshold $B = 0$. Conversely, when $L_2 = \infty$, p_2 covers the domain with $\mu_2 \geq 0$ and p_3 corresponds to $\mu_2 < 0$.

Tight performance bounds for the residual gradient in the nonsmooth case are conjectured in Section 4.2.3.

4.2.2 On the tightness of residual gradient bounds

Proposition 4.2.1 (Tightness of regimes p_1, p_2)**.** *Regimes p_1 and p_2, in Theorem 4.2.1 are tight for any number of iterations N, i.e., the corresponding lower bounds on the objective decrease cannot be improved.*

In Section 4.A we provide worst-case examples proving Proposition 4.2.1 in the case where both functions are smooth. The tightness of regime p_2, for the specific decomposition $f_1 \in \mathcal{F}_{0,L_1}$ and $f_2 \in \mathcal{F}_{0,\infty}$, is demonstrated by a worst-case construction in [3, Example 3.1]; using symmetry arguments, this example can be adapted to show the tightness of regime p_1 when f_1 is nonsmooth.

Conjecture 4.2.1 (Tightness of regime p_3)**.** *The DCA rate corresponding to regime p_3 from Corollary 4.2.1 is tight for any number of iterations N.*

Conjecture 4.2.2 (Tightness of regimes p_4, p_5, p_6)**.** *Regimes p_4, p_5, p_6 in Theorem 4.2.1 are tight after one iteration, i.e., the corresponding lower bounds on the objective decrease after one step cannot be improved.*

Proposition 4.2.1, Conjecture 4.2.1 and Conjecture 4.2.2 assert that our set of six inequalities from Theorem 4.2.1 represents the tightest possible characterization after a single iteration, in terms of the residual gradient. Specifically, there exist (separate) function examples for which each of these inequalities holds with equality.

Our numerical investigations show that the sublinear rates for p_4, p_5, p_6 in Corollary 4.2.1 are not tight beyond a single iteration. In the standard case of DCA with F being nonconvex and nonconcave ($\mu_1, \mu_2 < \min\{L_1, L_2\}$), the threshold condition $B = 0$ separates the tight and non-tight regimes in Corollary 4.2.1.

Conjectures 4.2.1 and 4.2.2 are supported by extensive numerical evidence obtained through the solution of a large number of performance estimation problems, spanning the range of parameters L_1, L_2, μ_1, μ_2; we illustrate this evidence in Section 4.A. Additionally, Figure 4.2.1, supporting the tightness of our bounds after one iteration, is generated by fixing μ_1 and L_1, creating a sampling grid for μ_2 and L_2, and solving the instances of (4.11) for $N = 1$.

The results align with the *exact* values of $p_1, \ldots, p_6$ in their respective regions. Conjecture 4.2.1, regarding tightness of regime p_3 after N iterations, was established through a similar sampling grid with $N \leq 20$, confirming that it yields tight sublinear rates.

4.2.3 Tight bounds in residual gradient for any number of iterations

Within this section we formulate the tight rates in the domains of regimes p_4, p_5, p_6, whose one-iteration analysis does not extend tightly for arbitrary number of steps. These conjectures are inspired by the connection with gradient descent (see Section 4.6 below) and have a direct correspondence with the theorems given in Chapter 3. GD is a particular case of the PGD, which is iteration equivalent to DCA when f_2 is smooth.

Conjecture 4.2.3 describes the exact behaviour when F is nonconvex-nonconcave, i.e., $L_2 > \mu_1$ and $L_1 > \mu_2$, while Conjecture 4.2.4 describes the exact behaviour when F is (strongly) convex, i.e., $L_2 \leq \mu_1$ and $L_1 > \mu_2$. Together, their domains cover that of p_5. In Section 4.A we show some numerical evidence that our conjectured expressions match the PEP numerical certificates.

Conjecture 4.2.3 extends Theorem 3.2.3 from Chapter 3, which is formulated for smooth nonconvex objectives. Conjecture 4.2.4 extends Theorem 3.2.2 and Theorem 3.2.1 from Chapter 3, which are formulated for smooth strongly convex and convex objective functions, respectively.

In all our rates within this section, to remove the dependence of the initial condition on the last iterate, the initial objective gap can be further upper bounded by $F(x^0) - F_{lo}$.

Conjecture 4.2.3 (Tight DCA rates for the residual gradient for general nonconvex-nonconcave F)**.** *Let* $\mu_2 < 0 < \mu_1 < L_2$, $\mu_1 + \mu_2 > 0$, $\mu_1^{-1} + \mu_2^{-1} + L_2^{-1} < 0$ *and* $N \geq 2$. *Then*

$$\frac{1}{2\mu_1} \min_{0 \leq k \leq N} \left\{ \|g_1^k - g_2^k\|^2 \right\} \leq \frac{F(x^0) - F(x^N)}{\min\left\{ P_N\left(\frac{L_2}{\mu_1}, \frac{\mu_2}{\mu_1}\right), \; E_N\left(\frac{\mu_2}{\mu_1}\right) \right\}},$$

where $P_N(\eta, \rho) := \frac{(1+\eta)(1+\rho)}{\eta+\rho}\left[N + \frac{(1-\eta)(1-\rho)}{\eta-\rho} \sum_{k=1}^{N} \left[E_N(\eta) - E_N(\rho) \right]_+ \right]$.

Conjecture 4.2.4 (Tight DCA rates for the residual gradient for general (strongly) convex F)**.** *Let* $\mu_1 > 0$, $\mu_2 \in \mathbb{R}$, *with* $\mu_1 + \mu_2 > 0$, $L_2 \in (0, \mu_1]$,

$N \geq 1$. *Then*

$$\frac{1}{2\mu_1} \min_{0\leq k\leq N} \left\{\|g_1^k - g_2^k\|^2\right\} \leq \frac{F(x^0) - F(x^N)}{\min\left\{E_N\left(\frac{L_2}{\mu_1}\right),\ E_N\left(\frac{\mu_2}{\mu_1}\right)\right\}}.$$

One can always upper bound the r.h.s. using that $F(x^N) \geq F_{lo}$. Remarkably, neither rate exhibits an explicit dependence on $L_1 > \mu_1 > 0$; hence, both apply to the smooth case $L_1 < \infty$ and the nonsmooth case $L_1 = \infty$. However, we recall that we implicitly assume exact oracle access to a subgradient of the conjugate f_1^*, which, when $L_1 < \infty$ corresponds to smooth functions f_1, typically entails lower computational effort than in the nonsmooth case. Additionally, the dependence appears only through the *ratios* of the remained parameters, namely $\frac{L_2}{\mu_1}$ and $\frac{\mu_2}{\mu_1}$. Anticipating the results for PGD in Section 4.6 (including the gradient descent case), these ratios correspond exactly to the factors $1-\gamma L$ and $1-\gamma\mu$, respectively, which determine the performance bounds for gradient descent in Chapter 3 (see Theorems 3.2.1 to 3.2.3 therein).

On subdomains of Conjectures 4.2.3 and 4.2.4 we are able to provide proofs using distance-1 interpolation inequalities. For the rest, we observed from the PEP numerics that distance-2 conditions are necessary, but unlike in the particular gradient descent case, we were not able to formally demonstrate them.

Theorem 4.2.2 (Tight DCA rates for the residual gradient for (strongly) convex F in regime p_4)**.** *Let $f_1 \in \mathcal{F}_{\mu_1,L_1}$ and $f_2 \in \mathcal{F}_{\mu_2,L_2}$, with $L_2 \in [0,\mu_1]$, $\mu_1+\mu_2 > 0$ and $\mu_2 \in \left[\frac{-L_2\mu_1}{L_2+\mu_1}, L_2\right)$. Then after N iterations of DCA starting from x^0 it holds:*

$$\frac{1}{2\mu_1}\|g_1^N - g_2^{N+1}\|^2 \leq \frac{F(x^0) - F(x^N)}{E_N\left(\frac{L_2}{\mu_1}\right)},$$

where $g_1^N \in \partial f_1(x^N)$ and $g_2^N \in \partial f_2(x^N)$.

Theorem 4.2.3 (Tight DCA rates for the residual gradient for nonconvex-nonconcave F with $L_2+\mu_2 \leq 0$ (subdomain of regime p_5))**.** *Let $f_1 \in \mathcal{F}_{\mu_1,L_1}$ and $f_2 \in \mathcal{F}_{\mu_2,L_2}$, with $\mu_1+\mu_2 > 0$ and $L_2+\mu_2 \leq 0$. Then after N iterations of DCA starting from x^0 it holds*

$$\frac{1}{2\mu_1}\|g_1^N - g_2^{N+1}\|^2 \leq \frac{F(x^0) - F(x^N)}{E_N\left(\frac{\mu_2}{\mu_1}\right)},$$

where $g_1^N \in \partial f_1(x^N)$ and $g_2^N \in \partial f_2(x^N)$.

Theorems 4.2.2 and 4.2.3 are proved later in Section 4.5.1, where certain descent lemmas are telescoped.

For completeness of exposure, Conjecture 4.2.5 below covers the case of minimizing a concave objective F and shows how fast the divergence occurs in the worst-case.

Conjecture 4.2.5 (Tight DCA rates for the residual gradient for (strongly) concave F (regime p_6))**.** *Let $0 \leq \mu_1 < L_1 \leq \mu_2 < L_2$, $N \geq 1$. Then*

$$\frac{1}{2\mu_2} \min_{0 \leq k \leq N} \left\{ \|g_1^k - g_2^k\|^2 \right\} \leq \frac{F(x^0) - F(x^N)}{\min\left\{ E_N\left(\frac{L_1}{\mu_2}\right), \, E_N\left(\frac{\mu_1}{\mu_2}\right) \right\}}.$$

4.3 Performance bounds in gradient mapping

This section focuses on deriving and discussing performance bounds for the best gradient mapping criterion, assuming that at least one of the two functions is smooth. In Section 4.3.1, we present a single-iteration analysis that yields six regimes and corresponding to sublinear rates, some of which are not tight. Tight bounds for an arbitrary horizon, some of which remain conjectural, are provided in Section 4.3.2.

4.3.1 Sublinear convergence rates in gradient mapping

For all results in this section, the gap $F(x^0) - F(x^{N+1})$ can be upper bounded by $F(x^0) - F_{\text{lo}}$.

Theorem 4.3.1 (DCA sublinear rates for the gradient mapping)**.** *Let $f_1 \in \mathcal{F}_{\mu_1, L_1}$ and $f_2 \in \mathcal{F}_{\mu_2, L_2}$ satisfy Assumptions 4.1.1 and 4.1.2, and assume $\mu_1 + \mu_2 > 0$. Then, after $N+1 \geq 2$ iterations of DCA starting from x^0, it holds that*

$$\tfrac{1}{2} \min_{0 \leq k \leq N} \{\|x^k - x^{k+1}\|^2\} \leq \frac{F(x^0) - F(x^{N+1})}{(\mu_1 + \mu_2) + r_i(\mu_1, L_1, \mu_2, L_2)N}, \tag{4.4}$$

with $i = 1, \ldots, 6$, and the expressions for the regimes r_i are given in Table 4.3.1.

The proof is deferred to Section 4.5.2. All regimes are exact for $N = \{0, 1\}$ and their domains span the entire range of parameters satisfying the condition $\mu_1 + \mu_2 > 0$, which is sufficient for decreasing the objective (see Proposition 4.1.2).

Table 4.3.1: Coefficients r_i from Theorem 4.3.1 characterizing the performance in gradient mapping after one iteration. Their domains cover all parameter space under the condition $\mu_1 + \mu_2 > 0$. Regimes $p_{\{1,2,3\}}$ are tight for any number of iterations. (s.c.: strongly convex; n.c.: nonconvex; ccv.: concave; n.ccv.: nonconcave)

Regime	Domain		Description
$r_1 = \mu_1 + \mu_2 + \frac{(\mu_1 - \mu_2)^2}{L_2 - \mu_2}$	$\mu_1 \geq \mu_2 \geq 0$	$L_2 > \mu_1$	f_1, f_2 convex F n.c.-n.ccv.
$r_2 = \mu_1 + \mu_2 + \frac{(\mu_1 - \mu_2)^2}{L_1 - \mu_1}$	$\mu_2 \geq \mu_1 \geq 0$	$L_1 > \mu_2$	f_1, f_2 convex F n.c.-n.ccv.
$r_3 = \frac{(\mu_1 + \mu_2)(L_2 + \mu_1)}{L_2 + \mu_2}$	$\mu_2 \in [\frac{-L_2\mu_1}{L_2 + \mu_1}, 0)$	$L_2 > \mu_1 > 0$	f_1 s.c., f_2 n.c. F n.c.-n.ccv.
$r_4 = \frac{\mu_1^2(L_2 + \mu_1)}{L_2^2}$	$\mu_2 \in \big[\frac{-L_2\mu_1}{L_2 + \mu_1}, L_2\big)$	$L_2 \in [0, \mu_1]$	f_1 s.c., f_2 n.c. F convex
$r_5 = \frac{\mu_1^2(\mu_1 + \mu_2)}{\mu_2^2}$	$\mu_2 < 0$, $\mu_2 \in \big(-\mu_1, \frac{-L_2\mu_1}{L_2 + \mu_1}\big]$		f_1 s.c., f_2 n.c. F n.ccv.
$r_6 = \frac{\mu_2^2(L_1 + \mu_2)}{L_1^2}$	$\mu_2 \geq \mu_1 \geq 0$	$L_1 \in (0, \mu_2]$	f_1 convex, f_2 s.c. F ccv.

Based on extensive numerical experiments using PEP, we determine that regimes $p_{\{1,2,3\}}$ are tight for any number of iterations. Regimes r_1 and r_2 are derived in [3, Proposition 3.1], by applying Toland duality [122] on the rates measuring the residual gradient criterion. Theorem 4.3.1 gives more explicit expressions and includes in the proofs a simpler derivation based on Lemmas 4.5.4 and 4.5.5. These regimes assume F is nonconvex-nonconcave ($L_1 > \mu_2$ and $L_2 > \mu_1$), with the active one indexed by $\arg\max_i\{\mu_1, \mu_2\}$. All other regimes are new. Figure 4.3.1 shows the domains of all possible six regimes holding after two iterations.

When F is nonconvex-nonconcave, regime r_3 is active when f_2 is weakly-convex with $\mu_2 \in [\frac{-L_2\mu_1}{L_2+\mu_1}, 0)$. When μ_2 is even lower, then regime r_5 is active. This is however not tight for $N \geq 2$, where a more complicated analysis is required. The exact rates expressions in this case are given in Conjecture 4.3.1.

When F is (strongly) convex ($\mu_1 \geq L_2$), the sublinear rates of regimes r_4 and r_5 are tight only for $N = 1$. These regimes are separated in this case by the threshold $\mu_2 = \frac{-L_2\mu_1}{L_2+\mu_1}$. When $N \geq 2$, we prove a linear rate for regime r_4 in Theorem 4.3.2, which is confirmed to be exact using PEP. Within the particular range of r_5 satisfying the condition $L_2 + \mu_2 \leq 0$, we demonstrate a linear rate in Theorem 4.3.3. For the rest domain of r_5, we conjecture the tight rates in Conjecture 4.3.1. Finally, regime r_6 characterizes the iterates over *concave*

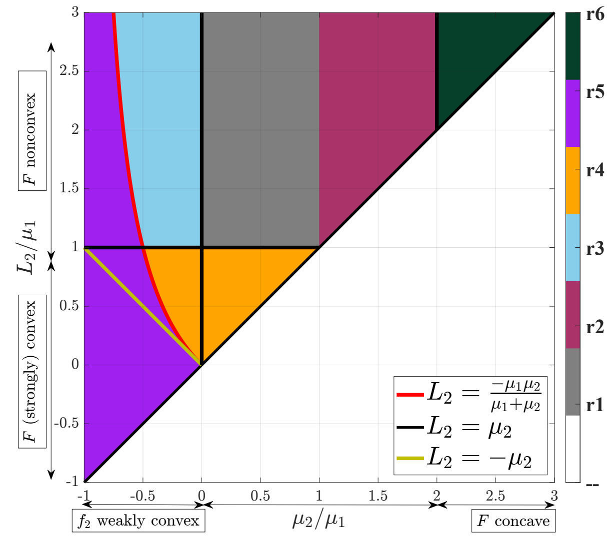

Figure 4.3.1: Domains of all tight regimes after two DCA iterations as proved in Theorem 4.3.1, given $\frac{L_1}{\mu_1} = 2$ and $\mu_1 > 0$. Regimes r_1, r_2 and r_3 are tight for any number of iterations $N \geq 2$. For the domain of r_4 we proved a linear rate for $N \geq 3$ (Theorem 4.3.2); same for the subdomain of r_5 satisfying the condition $L_2 + \mu_2 \leq 0$ (Theorem 4.3.3). For r_6 and the rest of domain of r_5 we conjecture rates in Section 4.3.2.

objectives F and is included for a complete covering of the parameter space.

One nonsmooth term

Corollary 4.3.1 (DCA sublinear rates for the gradient mapping with one nonsmooth term)**.** *Let $f_1 \in \mathcal{F}_{\mu_1,L_1}$ and $f_2 \in \mathcal{F}_{\mu_2,L_2}$, where exactly one function f_1 or f_2 is smooth, and assume that $\mu_1 + \mu_2 > 0$. Then, after $N + 1 \geq 2$ iterations of* (DCA) *starting from x^0, it holds that*

$$\tfrac{1}{2} \min_{0\leq k\leq N} \{\|x^k - x^{k+1}\|^2\} \leq \frac{F(x^0) - F(x^{N+1})}{(\mu_1 + \mu_2) + r_i(L_1, L_2, \mu_1, \mu_2)N},$$

where the regimes r_i are as follows: for $L_1 = \infty$, the regimes r_i with $i = \{1, 2, 3, 4, 5\}$ are given in Table 4.3.2; for $L_2 = \infty$, the regimes r_i with $i = \{1, 2, 3, 6\}$ are given in Table 4.3.3.

When f_1 is nonsmooth, regime r_6 disappears (as F cannot be concave in this case), while the expression of regime r_2 simplifies to the standard decrease

$\mu_1 + \mu_2$ given in Proposition 4.1.2. Remarkably, the expressions of regimes r_3, r_4 and r_5 are untouched in the limit $L_1 = \infty$.

Recall that, when f_2 is nonsmooth, the critical points are not necessarily the stationary ones. In this case, the progress in the area of regimes r_1 and r_3 reduces to the standard decrease from Proposition 4.1.2, while regimes r_2 and r_6 are untouched.

The fact that in the nonsmooth case some expressions reduce to the standard decrease $\mu_1 + \mu_2$ is due to the fact that the algorithm cannot exploit anymore the additional smoothness information. Moreover, if both functions would be nonsmooth, then this decrease is the maximum that can be proved on the iteration progress. In this case, we provide an alternative metric in Section 4.4.

Table 4.3.2: Performance in gradient mapping when f_1 is nonsmooth ($L_1 = \infty$), through coefficients r_i from Corollary 4.3.1. Their domains cover all parameter space under the condition $\mu_1 + \mu_2 > 0$. (s.c.: strongly convex; n.c.: nonconvex; ccv.: concave; n.ccv.: nonconcave)

Regime	Domain		Description
$r_1 = \mu_1 + \mu_2 + \frac{(\mu_1-\mu_2)^2}{L_2-\mu_2}$	$\mu_1 \geq \mu_2 \geq 0$	$L_2 > \mu_1$	f_1, f_2 convex F n.c.-n.ccv.
$r_2 = \mu_1 + \mu_2$	$\mu_2 \geq \mu_1 \geq 0$	$L_1 > \mu_2$	f_1, f_2 convex F n.c.-n.ccv.
$r_3 = \frac{(\mu_1+\mu_2)(L_2+\mu_1)}{L_2+\mu_2}$	$\mu_2 \in [\frac{-L_2\mu_1}{L_2+\mu_1}, 0)$	$L_2 > \mu_1 > 0$	f_1 s.c., f_2 n.c. F n.c.-n.ccv.
$r_4 = \frac{\mu_1^2(L_2+\mu_1)}{L_2^2}$	$\mu_2 \in \left[\frac{-L_2\mu_1}{L_2+\mu_1}, L_2\right)$	$L_2 \in [0, \mu_1]$	f_1 s.c., f_2 n.c. F convex
$r_5 = \frac{\mu_1^2(\mu_1+\mu_2)}{\mu_2^2}$	$\mu_2 < 0$, $\mu_2 \in \left(-\mu_1, \frac{-L_2\mu_1}{L_2+\mu_1}\right]$		f_1 s.c., f_2 n.c. F n.ccv.

4.3.2 Tight bounds in gradient mapping for any number of iterations

Within this section we formulate the tight rates for any number of iterations for the regimes where the analysis for $N = 1$, leading to sublinear rates, does not extend exactly for $N > 1$. We distinguish between two cases: (i) proofs using distance-1 inequalities, (ii) conjectured parts, whose proofs require using distance-2 inequalities. In Conjecture 4.3.1 we provide the tight rates (as identified using PEP) for all regimes requiring a different analysis from the one in the previous section.

Table 4.3.3: Performance in gradient mapping when f_2 is nonsmooth ($L_2 = \infty$), through coefficients r_i from Corollary 4.3.1. Their domains cover all parameter space under the condition $\mu_1 + \mu_2 > 0$. (s.c.: strongly convex; n.c.: nonconvex; ccv.: concave; n.ccv.: nonconcave)

Regime	Domain		Description
$r_1 = r_3 = \mu_1 + \mu_2$	$\mu_1 \geq \mu_2 \geq 0$		f_1, f_2 convex F n.c.-n.ccv.
$r_2 = \mu_1 + \mu_2 + \frac{(\mu_1-\mu_2)^2}{L_1-\mu_1}$	$\mu_2 \geq \mu_1 \geq 0$	$L_1 > \mu_2$	f_1, f_2 convex F n.c.-n.ccv.
$r_6 = \frac{\mu_2^2(L_1+\mu_2)}{L_1^2}$	$\mu_2 \geq \mu_1 \geq 0$	$L_1 \in (0, \mu_2]$	f_1 convex, f_2 s.c. F ccv.

Conjecture 4.3.1 (Tight DCA rates for the gradient mapping for general F)**.** *Let $f_1 \in \mathcal{F}_{\mu_1,L_1}$ and $f_2 \in \mathcal{F}_{\mu_2,L_2}$, with $\mu_1 \geq 0$ and $\mu_1 + \mu_2 > 0$. Consider $N+1 \geq 2$ iterations of DCA starting from x^0.*

i) If $\mu_2 < 0 < \mu_1 < L_2$, thus F is nonconvex-nonconcave, then:

$$\frac{1}{2\mu_1} \min_{0\leq k\leq N} \left\{\|\mu_1(x^k - x^{k+1})\|^2\right\} \leq \frac{F(x^0) - F(x^{N+1})}{1 + \frac{\mu_2}{\mu_1} + \min\left\{P_N\left(\frac{L_2}{\mu_1}, \frac{\mu_2}{\mu_1}\right), E_N\left(\frac{\mu_2}{\mu_1}\right)\right\}}, \tag{4.5}$$

where $P_N(\eta,\rho) := \frac{(1+\eta)(1+\rho)}{\eta+\rho}\left(N + \frac{(1-\eta)(1-\rho)}{\eta-\rho}\sum_{k=0}^{N}[E_k(\eta) - E_k(\rho)]_+\right)$.

ii) If $L_2 \leq \mu_1$ and $\mu_1 > 0$, thus F is (strongly) convex, then:

$$\frac{1}{2\mu_1}\|\mu_1(x^k - x^{k+1})\|^2 \leq \frac{F(x^0) - F(x^{N+1})}{1 + \frac{\mu_2}{\mu_1} + \min\left\{E_N\left(\frac{L_2}{\mu_1}\right), E_N\left(\frac{\mu_2}{\mu_1}\right)\right\}}. \tag{4.6}$$

iii) If $0 \leq \mu_1 < L_1 \leq \mu_2 < L_2$, thus F is (strongly) concave, then:

$$\frac{1}{2\mu_2}\|\mu_2(x^k - x^{k+1})\|^2 \leq \frac{F(x^0) - F(x^{N+1})}{1 + \frac{\mu_1}{\mu_2} + \min\left\{E_N\left(\frac{L_1}{\mu_2}\right), E_N\left(\frac{\mu_1}{\mu_2}\right)\right\}}. \tag{4.7}$$

For $N = 1$, Conjecture 4.3.1 includes the rates corresponding to regimes r_3 and r_5 in (4.5) when F is nonconvex, of regimes r_4 and r_5 in (4.6) when F is (strongly) convex and of regime r_6 in (4.7) when F is concave. The rate in (4.6) includes the ones proved in Theorem 4.3.2 (within the domain of regime r_4) and in Theorem 4.3.3. The expressions of the rates (4.5) and (4.6) are independent of L_1, hence they also hold when f_1 is ***nonsmooth***, i.e., $L_1 = \infty$. Note that the normalization of the iterates differences by μ_1 or μ_2 results in a homogeneous expression of the right-hand side denominators in the ratios $\frac{L_2}{\mu_1}$ and $\frac{\mu_2}{\mu_1}$, or $\frac{L_1}{\mu_2}$ and $\frac{\mu_1}{\mu_2}$, respectively.

Conjecture 4.3.1 is based on the exact rates developed for gradient descent derived in [100] (presented in Chapter 3) and is confirmed by PEP numerical experiments.

Observe that the conjectures for any number of iterations $N \geq 2$, stated either for residual gradient or gradient mapping, share the same r.h.s. under the right scaling. This is expected as they are both inspired by gradient descent where the two measures coincide.

Part of Conjecture 4.3.1 can be proved by using distance-1 interpolation inequalities, similarly to what we are able to prove in Section 4.2.3. The rest requires a more intricate analysis using distance-2 interpolation conditions, as shown by PEP numerical simulations.

Theorem 4.3.2 provides the exact rate in the area of regime r_4, whereas Theorem 4.3.3 gives the exact behaviour in the subdomain of r_5, bounded by the condition $L_2 + \mu_2 \leq 0$. In both cases, the proofs only involve combining inequalities on consecutive iterations given in Lemma 4.6.1.

Theorem 4.3.2 (Tight DCA rates for the gradient mapping for (strongly) convex F in regime r_4)**.** *Let $f_1 \in \mathcal{F}_{\mu_1,L_1}$ and $f_2 \in \mathcal{F}_{\mu_2,L_2}$, with $L_2 \in [0,\mu_1]$, $\mu_1+\mu_2 > 0$ and $\mu_2 \in \left[\frac{-L_2\mu_1}{L_2+\mu_1}, L_2\right)$. Then after $N+1 \geq 2$ iterations of DCA starting from x^0 it holds:*

$$\frac{1}{2\mu_1}\|\mu_1(x^N - x^{N+1})\|^2 \leq \frac{F(x^0)-F(x^{N+1})}{1+\frac{\mu_2}{\mu_1}+E_N\left(\frac{L_2}{\mu_1}\right)}.$$

Theorem 4.3.3 (Tight DCA rates for the gradient mapping for nonconvex-non-concave F with $L_2+\mu_2 \leq 0$ (subdomain of regime r_5))**.** *Let $f_1 \in \mathcal{F}_{\mu_1,L_1}$ and $f_2 \in \mathcal{F}_{\mu_2,L_2}$, with $\mu_1+\mu_2 > 0$ and $L_2+\mu_2 \leq 0$. Then after $N+1 \geq 2$ iterations of DCA starting from x^0 it holds*

$$\frac{1}{2\mu_1}\|\mu_1(x^N - x^{N+1})\|^2 \leq \frac{F(x^0)-F(x^{N+1})}{1+\frac{\mu_2}{\mu_1}+E_N\left(\frac{\mu_2}{\mu_1}\right)}.$$

Theorems 4.3.2 and 4.3.3 are proved later in Section 4.5.2.

Remark 4.3.1 (Parallel between residual gradient and gradient mapping regimes)**.** *Regimes r_1 and r_2 can be derived from regimes p_1 and p_2 via Toland* *duality [122], as shown in [3, Proposition 3.1] (after a minor reorganization of the presentation to make the regimes explicit). Specifically, r_1 and r_2 can be obtained from p_1 and p_2 by the substitutions*

$$L_1 \mapsto \mu_2^{-1}, \qquad L_2 \mapsto \mu_1^{-1}, \qquad \mu_1 \mapsto L_2^{-1}, \qquad \mu_2 \mapsto L_1^{-1},$$

where the l.h.s. denotes the parameters for residual gradient regimes p_i and the r.h.s. denotes those in the gradient mapping regimes r_i. There is no clear correspondence between r_3 and p_3. Regimes p_4, p_5, and p_6 in the residual gradient analysis yield analogous rates in the gradient mapping setting, corresponding to regimes r_4, r_5, and r_6, respectively, under the substitution $G^{k\,/\,k+1}$ (residual gradient) by $\mu_1 \Delta x^{k\,/\,k+1}$ (gradient mapping).

4.4 Performance bounds with both nonsmooth

Within this section, we assume $L_1 = L_2 = \infty$. By taking the limit in rates from Corollary 4.2.2 (see Table 4.2.2), we observe that the progress after one iteration becomes $F(x) - F(x^+) \geq 0$ in all regimes, without any convergence guarantee. Therefore, we introduce the following metric, i.e., the Bregman divergence, to analyze the progress of (DCA):

$$\begin{aligned} T(x) &:= f_1(x) - f_2(x) - \inf_w \{f_1(w) - f_2(x) - \langle g_2\,,\, w - x\rangle\} \\ &= f_1(x) - f_1(x^+) - \langle g_1^+\,,\, x - x^+\rangle, \end{aligned} \tag{4.8}$$

where $g_2 \in \partial f_2(x)$ and $g_1^+ \in \partial f_1(x^+)$, the second identity resulting from the (DCA) iteration definition. This measure is used for studying the convergence of algorithms on nonconvex problems such as proximal gradient methods [64, Theorem 5] or the Frank-Wolfe algorithm [51, Equation (2.6)]. For DCA, the performance measured by the progress in $T(x)$ is derived in [3, §4] when both f_1 and f_2 are convex; the same result is obtained independently in [136, Corollary 3.2]. We extend their analysis to encompass weakly convex functions and provide a simpler and stronger proof in the convex case.

Since $\mu_1 \geq 0$, it holds $T(x) \geq \frac{1}{2}\mu_1\|x - x^+\|^2 \geq 0$; this comes from the subgradient inequality on f_1 and second line of (4.8). Furthermore, if $T(x^k) = 0$ is achieved for some iterate x^k, then $F(x^{k+1}) = F(x^k)$ and hence x^k is a critical point. The metric $T(x)$ is nothing but the Bregman divergence associated to f_1 at (x, x^+). In this view, the DCA is analyzed in [48] as being the Bregman proximal point algorithm.

Theorem 4.4.1 (DCA one-step descent for the Bregman divergence)**.** *Let $f_1 \in \mathcal{F}_{\mu_1,\infty}$ and $f_2 \in \mathcal{F}_{\mu_2,\infty}$, with $\mu_1 + \mu_2 > 0$ or $\mu_1 = \mu_2 = 0$. One iteration* *of* (DCA) *satisfies*

$$\tfrac{1}{2}\mu_2\|x - x^+\|^2 + T(x) \le F(x) - F(x^+), \qquad \textit{if } \mu_2 \ge 0; \tag{4.9}$$

$$\frac{\mu_1 + \mu_2}{\mu_1} T(x) \le F(x) - F(x^+), \qquad \textit{if } \mu_2 < 0. \tag{4.10}$$

The proof of Theorem 4.4.1 is given in Section 4.5.3.

Theorem 4.4.2 (DCA sublinear rates for the Bregman divergence)**.** *Let* $f_1 \in \mathcal{F}_{\mu_1,\infty}$ *and* $f_2 \in \mathcal{F}_{\mu_2,\infty}$*, with* $\mu_1 + \mu_2 > 0$ *or* $\mu_1 = \mu_2 = 0$*. Consider* N *iterations of* (DCA) *starting from* x^0*. Then, if* $\mu_2 \ge 0$*:*

$$\min_{0\le k\le N-1}\{T(x^k)\} + \frac{\mu_2}{2}\min_{0\le k\le N-1}\{\|x^k - x^{k+1}\|^2\} \le \frac{F(x^0) - F_{lo}}{N},$$

whereas if $\mu_2 < 0$ *we have:*

$$\min_{0\le k\le N-1}\{T(x^k)\} \le \frac{\mu_1}{\mu_1 + \mu_2}\frac{F(x^0) - F_{lo}}{N}.$$

Proof. Telescoping the inequalities from Theorem 4.4.1, taking the minimum over all $T(x^k)$, and using the trivial bound $F(x^N) \ge F_{\text{lo}}$concludes the proof. □

When $\mu_2 \ge 0$, the previous work in [3] only includes the first term $\min_{0\le k\le N-1}\{T(x^k)\}$ appearing in the left-hand side of the first inequality above, while the same result as ours is stated in a more general and compact form using Bregman divergence in [48, Theorem 5].

4.5 Proofs of performance bounds

In this section we present the proofs for the upper bounds derived in the previous section, which are inspired by and confirmed using the performance estimation methodology. The following general performance estimation setup for DCA can be easily integrated into one of the specialized software packages PESTO (in MATLAB; Taylor et al. [115]) or PEPit (in Python; Goujaud et al. [56]):

$$\begin{aligned} &\text{maximize} \quad \frac{\tfrac{1}{2}\min\limits_{0\le k\le N}\{\|g_1^k - g_2^k\|^2\}}{F(x^0) - F(x^N)} \\ &\text{subject to} \quad \{(x^k, g_{1,2}^k, f_{1,2}^k)\}_{k\in\mathcal{I}} \text{ satisfy (2.7)} \\ &\qquad\qquad\qquad g_1^{k+1} = g_2^k \quad k \in \{0,\dots,N-1\}. \end{aligned} \tag{4.11}$$

The decision variables are x^k, g_1^k, g_2^k, f_1^k, f_2^k, with $k \in \mathcal{I} = \{0, \ldots, N\}$. A comprehensive formulation of the PEP for DCA can be found in [3]. The numerical solutions of (4.11) aided our derivations, enabling verification of the rates from Corollary 4.2.1, conjecturing their tightness, and selecting for the proofs the necessary interpolation inequalities from the complete set given in Theorem 2.2.1.

While employing PEP can sometimes make proofs difficult to comprehend, we offer demonstrations of the various regimes in Section 4.5.1 (for residual gradient) and in Section 4.5.2 based on standard descent lemmas. This involves combining interpolation inequalities for consecutive iterations, with the proof for each of the six regimes requiring a distinct weighted combination of them.

Notation. For compactness of the proofs, at each iteration k we denote $\Delta x^k := x^k - x^{k+1}$, $f_1^k = f_1(x^k)$, $f_2^k = f_2(x^k)$, $\Delta F(x^k) := F(x^k) - F(x^{k+1})$, $g_1^k \in \partial f_1(x^k)$, $g_2^k \in \partial f_2(x^k)$, $G^k := g_1^k - g_2^k$. Recall that $g_1^{k+1} = g_2^k$ for all DCA iterations. We omit the explicit dependence on the curvature bounds μ_1, L_1, μ_2 and L_2.

The main theorems rely on a series of technical lemmas that provide the building blocks for our convergence analysis. These lemmas are organized by performance metric and proof technique.

Key ingredient for the proofs. Given $\mu < L$ and $f \in \mathcal{F}_{\mu,L}$, from Theorem 2.2.1 it holds $\forall x, y \in \mathbb{R}^d$, with $g^x \in \partial f(x)$, $g^y \in \partial f(y)$, that

$$f(x) - f(y) - \langle g^y, x - y\rangle \geq \frac{\mu}{2}\|x - y\|^2 + \frac{1}{2(L-\mu)}\|g^x - g^y - \mu(x - y)\|^2. \tag{4.12}$$

4.5.1 Proofs of bounds in residual gradient

We distinguish two types of supporting lemmas for residual gradient analysis:

- **Individual regime proofs**: The one-step descent inequalities for regimes p_1 through p_6 are proved in a single consolidated proof of Theorem 4.2.1 (see Section 4.5.1), rather than via separate regime-by-regime arguments.
- **Linear regime descent lemmas**: Lemma 4.5.1 provides the descent inequality for linear convergence in regime p_4, while Lemma 4.5.2 establishes the corresponding result for the subdomain of regime p_5 satisfying $L_2 + \mu_2 \leq 0$ (Section 4.5.1).

Table 4.5.1: Exact decrease in residual gradient after one step: $\Delta F(x^k) \geq \sigma_i \frac{1}{2}\|G^k\|^2 + \sigma_i^+ \frac{1}{2}\|G^{k+1}\|^2$ (see Theorem 4.2.1), where $\sigma_i, \sigma_i^+ \geq 0$ and $p_i = \sigma_i + \sigma_i^+$, with $i = 1, \ldots, 6$. Scalar $\alpha_i \geq 0$ is a parameter of the proofs.

Regime	σ_i	σ_i^+	α_i
p_1	$L_2^{-1}\dfrac{L_2-\mu_1}{L_1-\mu_1}$	$L_2^{-1}\left(1+\dfrac{L_2^{-1}-L_1^{-1}}{\mu_1^{-1}-L_1^{-1}}\right)$	$\dfrac{\mu_1}{L_2}\dfrac{L_1-L_2}{L_1-\mu_1}$
p_2	$L_1^{-1}\left(1+\dfrac{L_1^{-1}-L_2^{-1}}{\mu_2^{-1}-L_2^{-1}}\right)$	$L_1^{-1}\dfrac{L_1-\mu_2}{L_2-\mu_2}$	$\dfrac{\mu_2}{L_1}\dfrac{L_2-L_1}{L_2-\mu_2}$
p_3	$\dfrac{L_1^{-1}(\mu_1^{-1}+\mu_2^{-1}+L_2^{-1})}{\mu_1^{-1}+\mu_2^{-1}+L_2^{-1}-L_1^{-1}}$	$\dfrac{1}{L_2+\mu_2}$	$\dfrac{-\mu_2}{L_2+\mu_2}$
p_4	0	$\dfrac{L_2+\mu_1}{L_2^2}$	$\dfrac{\mu_1}{L_2}$
p_5	0	$\dfrac{\mu_1+\mu_2}{\mu_2^2}$	$\dfrac{\mu_1+\mu_2}{-\mu_2}$
p_6	$\dfrac{L_1+\mu_2}{L_1^2}$	0	$\dfrac{\mu_2}{L_1}$

Proof of one iteration descent (Theorem 4.2.1)

In Table 4.5.1 we summarize the coefficients for all regimes, along with the corresponding multipliers α_i, which are utilized in the subsequent proofs of each regime in part. Within this table we adopt the notation corresponding to $x^k \mapsto x^{k+1}$ instead of the former notation with $x \mapsto x^+$. Each row of this table can be read as a descent lemma in its own right, and all of the statements are proved together in a single proof, with each row addressed in a separate paragraph.

Proof of Theorem 4.2.1. By writing (4.12) for function f_1 with the iterates (x^k, x^{k+1}) we obtain:

$$f_1(x^k) - f_1(x^{k+1}) - \langle g_1^{k+1}, \Delta x^k \rangle \geq \frac{1}{2L_1}\|G^k\|^2 + \frac{\mu_1}{2L_1(L_1-\mu_1)}\|G^k - L_1 \Delta x^k\|^2 \tag{4.13}$$

and for function f_2 with the iterates (x^{k+1}, x^k) we get:

$$f_2(x^{k+1}) - f_2(x^k) + \langle g_2^k, \Delta x^k \rangle \geq \frac{1}{2L_2}\|G^{k+1}\|^2 + \frac{\mu_2}{2L_2(L_2-\mu_2)}\|G^{k+1} - L_2\Delta x^k\|^2. \tag{4.14}$$

Summing them up and performing simplifications we get:

$$\Delta F^k \geq \frac{1}{2L_1}\|G^k\|^2 + \frac{\mu_1}{2L_1(L_1-\mu_1)}\|G^k - L_1\Delta x^k\|^2 + \\ + \frac{1}{2L_2}\|G^{k+1}\|^2 + \frac{\mu_2}{2L_2(L_2-\mu_2)}\|G^{k+1} - L_2\Delta x^k\|^2. \tag{4.15}$$

By writing (4.12) for function f_2 with the iterates (x^k, x^{k+1}):

$$f_2(x^k) - f_2(x^{k+1}) - \langle g_2^{k+1}, \Delta x^k \rangle \geq \frac{1}{2L_2}\|G^{k+1}\|^2 + \frac{\mu_2}{2L_2(L_2-\mu_2)}\|G^{k+1} - L_2\Delta x^k\|^2$$

summing it with (4.14), and using $G^{k+1} = g_2^k - g_2^{k+1}$, we obtain:

$$\langle G^{k+1}, \Delta x^k \rangle \geq \frac{1}{L_2}\|G^{k+1}\|^2 + \frac{\mu_2}{L_2(L_2-\mu_2)}\|G^{k+1} - L_2\Delta x^k\|^2. \tag{4.16}$$

Similarly, by writing (4.12) for function f_1 with the iterates (x^{k+1}, x^k) and summing it up with (4.13) we get:

$$\langle G^k, \Delta x^k \rangle \geq \frac{1}{L_1}\|G^k\|^2 + \frac{\mu_1}{L_1(L_1-\mu_1)}\|G^k - L_1\Delta x^k\|^2. \tag{4.17}$$

The proofs only involve adjusting the right-hand side of (4.15) using either inequality (4.16) or inequality (4.17), weighted by scalars $\alpha > 0$, which depend on the curvatures (see Table 4.5.1). Specifically, to establish regimes $p_{1,3,4,5}$ we substitute α in:

$$\Delta F(x^k) \geq \frac{1}{2L_1}\|G^k\|^2 + \frac{\mu_1}{2L_1(L_1-\mu_1)}\|G^k - L_1\Delta x^k\|^2 + \\ \frac{1+2\alpha}{2L_2}\|G^{k+1}\|^2 + \frac{\mu_2(1+2\alpha)}{2L_2(L_2-\mu_2)}\|G^{k+1} - L_2\Delta x^k\|^2 - \alpha\langle G^{k+1}, \Delta x^k \rangle \tag{4.18}$$

and to demonstrate regimes $p_{2,6}$ we plug in α in:

$$\Delta F(x^k) \geq \frac{1}{2L_2}\|G^{k+1}\|^2 + \frac{\mu_2}{2L_2(L_2-\mu_2)}\|G^{k+1} - L_2\Delta x^k\|^2 + \frac{1+2\alpha}{2L_1}\|G^k\|^2 + \frac{\mu_1(1+2\alpha)}{2L_1(L_1-\mu_1)}\|G^k - L_1\Delta x^k\|^2 - \alpha\langle G^k, \Delta x^k\rangle.$$

Since the proof is entirely based on algebraic manipulations, by exploiting the symmetry on the right-hand side of the two inequalities we focus on demonstrating the regimes p_1, p_3, p_4, p_5. Regimes p_2 and p_6 are complementary to p_1 and p_4, respectively, under the condition $\mu_1 \geq 0$; specifically, $p_1 \iff p_2$ and $p_4 \iff p_6$ (see also Table 4.2.1). Their proofs can be obtained by interchanging: (i) the curvature indices 1 and 2; and (ii) G^k and G^{k+1}.

Regime p_1: It corresponds to $L_1 \geq L_2 > \mu_1 \geq 0$; in particular, $L_1 \geq \max\{L_2, \mu_1, \mu_2\}$. By setting $\alpha = \frac{\mu_1}{L_2}\frac{L_1-L_2}{L_1-\mu_1}$ in (4.18), after simplifications and completing the squares we get:

$$\Delta F(x^k) \geq \frac{L_2-\mu_1}{L_2(L_1-\mu_1)}\frac{\|G^k\|^2}{2} + \frac{\mu_1}{L_2(L_1-\mu_1)}\frac{\|G^k - L_2\Delta x^k\|^2}{2} + \frac{1}{L_2}\left[1 + \frac{\mu_1(L_1-L_2)}{L_2(L_1-\mu_1)}\right]\frac{\|G^{k+1}\|^2}{2} + \frac{\mu_1\frac{L_1}{L_2}\mu_2\left[\mu_1^{-1} + \mu_2^{-1} + L_2^{-1} - L_1^{-1}\left(2 + \frac{L_2}{\mu_2}\right)\right]}{(L_1-\mu_1)(L_2-\mu_2)}\frac{\|G^{k+1} - L_2\Delta x^k\|^2}{2}. \tag{4.19}$$

The weight of $\|G^k - L_2\Delta x^k\|^2$ is positive ($\mu_1 \geq 0$), while the weight of $\|G^{k+1} - L_2\Delta x^k\|^2$ is positive if either (i) $\mu_2 \geq 0$ or (ii) $\mu_2 < 0$ and $\mu_1^{-1} + \mu_2^{-1} + L_2^{-1} - L_1^{-1} \leq L_1^{-1}(1 + \frac{L_2}{\mu_2})$. Under these two cases, by neglecting both mixed terms we get:

$$\Delta F(x^k) \geq \frac{L_2-\mu_1}{L_2(L_1-\mu_1)}\frac{\|G^k\|^2}{2} + \frac{1}{L_2}\left(1 + \frac{L_2^{-1} - L_1^{-1}}{\mu_1^{-1} - L_1^{-1}}\right)\frac{\|G^{k+1}\|^2}{2},$$

with equality only if $G^k = G^{k+1} = L_2\Delta x^k$.

Regime p_3: It corresponds to $L_2 \geq \mu_1 \geq 0$ and negativity of the weight of $\|G^{k+1} - L_2\Delta x^k\|^2$ from the proof of p_1, i.e., $\mu_2 < 0$ and $\mu_1^{-1} + \mu_2^{-1} + L_2^{-1} - L_1^{-1} > L_1^{-1}(1 + \frac{L_2}{\mu_2})$. Since $L_2 \geq \mu_1 > -\mu_2$, we have $L_2 + \mu_2 > 0$. By setting $\alpha = \frac{-\mu_2}{L_2+\mu_2}$ in (4.18), after simplifications and completing a square including all weight of

$L_2\Delta x$ we get:

$$\Delta F(x^k) \geq \frac{L_1^{-1}(\mu_1^{-1}+\mu_2^{-1}+L_2^{-1})}{\mu_1^{-1}+\mu_2^{-1}+L_2^{-1}-L_1^{-1}}\frac{\|G^k\|^2}{2}+\frac{1}{L_2+\mu_2}\frac{\|G^{k+1}\|^2}{2}+$$
$$+\frac{\mu_1 L_1\mu_2(\mu_1^{-1}+\mu_2^{-1}+L_2^{-1}-L_1^{-1})}{L_2(L_1-\mu_1)(L_2-\mu_2)}\frac{\left\|\frac{1}{\frac{L_1}{L_2}+\frac{\mu_2(L_1-\mu_1)}{\mu_1(L_2+\mu_2)}}G^k-L_2\Delta x^k\right\|^2}{2}.$$

Since $\mu_2<0$, the weight of the mixed term is nonnegative only if $\mu_1^{-1}+\mu_2^{-1}+L_2^{-1}-L_1^{-1}\leq 0$. Further on, this implies that the weight of $\|G\|^2$ is nonnegative only if $\mu_1^{-1}+\mu_2^{-1}+L_2^{-1}\leq 0$, which is exactly the threshold condition $B\leq 0$, with B defined in (4.3). Finally, by neglecting the mixed term we get:

$$\Delta F(x^k) \geq \frac{L_1^{-1}(\mu_1^{-1}+\mu_2^{-1}+L_2^{-1})}{\mu_1^{-1}+\mu_2^{-1}+L_2^{-1}-L_1^{-1}}\frac{\|G^k\|^2}{2}+\frac{1}{L_2+\mu_2}\frac{\|G^{k+1}\|^2}{2},$$

with equality only if $G^k=L_2\left[\frac{L_1}{L_2}+\frac{\mu_2(L_1-\mu_1)}{\mu_1(L_2+\mu_2)}\right]\Delta x^k$.

Regime p_5**:** When $\mu_2<0$, if $\mu_1^{-1}+\mu_2^{-1}+L_2^{-1}>0$, the proof for regime p_3 breaks. Then we set $\alpha=\frac{\mu_1+\mu_2}{-\mu_2}>0$ in (4.18). After simplifications and completing the squares we get:

$$\begin{aligned}\Delta F(x^k) \geq \frac{\mu_1+\mu_2}{\mu_2^2}\frac{\|G^{k+1}\|^2}{2}+\frac{1}{L_1-\mu_1}\frac{\|G^k-\mu_1\Delta x^k\|^2}{2}+\\ \frac{\mu_1 L_2}{-\mu_2}\frac{\mu_1^{-1}+\mu_2^{-1}+L_2^{-1}}{L_2-\mu_2}\frac{\|G^{k+1}-\mu_2\Delta x^k\|^2}{2}.\end{aligned}\tag{4.20}$$

The second mixed term has a positive weight if $L_2(\mu_1^{-1}+\mu_2^{-1}+L_2^{-1})>0$. Note that this condition allows $L_2\leq 0$, bounded bellow through the necessary condition $L_2>\mu_2>-\mu_1$. After disregarding the mixed squares with positive weights, we obtain:

$$\Delta F(x^k)\geq\frac{\mu_1+\mu_2}{\mu_2^2}\frac{\|G^{k+1}\|^2}{2},$$

which holds with equality only if $G^k=\mu_1\Delta x^k$ and $G^{k+1}=\mu_2\Delta x^k$.

Regime p_4**:** Assume $0<L_2\leq\mu_1$, i.e., F is (strongly) convex, and $\mu_2(\mu_1^{-1}+\mu_2^{-1}+L_2^{-1})\geq 0$, where either $\mu_2\geq 0$ or $\mu_1>-\mu_2>0$. By setting $\alpha=\frac{\mu_1}{L_2}$ in (4.18), after simplifications and completing the squares we get:

$$\begin{aligned}\Delta F(x^k) \geq \frac{\mu_1+L_2}{L_2^2}\frac{\|G^{k+1}\|^2}{2}+\frac{1}{L_1-\mu_1}\frac{\|G^k-\mu_1\Delta x^k\|^2}{2}+\\ +\frac{\mu_1\mu_2}{L_2}\frac{\mu_1^{-1}+\mu_2^{-1}+L_2^{-1}}{L_2-\mu_2}\frac{\|G^{k+1}-L_2\Delta x^k\|^2}{2},\end{aligned}\tag{4.21}$$

where the mixed term can be disregarded to obtain:

$$\Delta F(x^k) \geq \frac{L_2 + \mu_1}{L_2^2} \frac{\|G^{k+1}\|^2}{2},$$

which holds with equality only if $G^k = \mu_1 \Delta x^k$ and $G^{k+1} = L_2 \Delta x^k$. □

Remark 4.5.1 (Generating multipliers α for single-iteration residual gradient descent)**.** *In the r.h.s. of inequality* (4.18) *we eliminated* Δx^k *by incorporating it in square including* G^k *and* G^{k+1}*, with the idea of maximizing the sum of weights of* $\|G^k\|$ *and* $\|G^{k+1}\|$*.*

Proof of sublinear rates (Corollary 4.2.1)

Proof of Corollary 4.2.1. From Theorem 4.2.1, by taking the minimum between the (sub)gradients residual norms in (4.1) we get:

$$\begin{aligned} F(x) - F(x^+) &\geq \sigma_i \tfrac{1}{2}\|g_1 - g_2\|^2 + \sigma_i^+ \tfrac{1}{2}\|g_1^+ - g_2^+\|^2 \\ &\geq p_i \tfrac{1}{2} \min\{\|g_1 - g_2\|^2, \|g_1^+ - g_2^+\|^2\}, \end{aligned}$$

where $p_i = \sigma_i + \sigma_i^+$, for $i = 1, \ldots, 6$, are given in Table 4.2.1. The rate (4.2) results by telescoping the above inequality for N iterations and taking the minimum among all (sub)gradients residual norms. A rate with respect to F_{lo} is obtained either by applying the trivial bound $F(x^0) - F(x^N) \geq F(x^0) - F_{lo}$ or, if $L_1 > \mu_2$, by using a tighter bound like the one demonstrated by Abbaszadehpeivasti et al. [3, Lemma 2.1]:

$$F(x^N) - F_{lo} \geq \frac{1}{2(L_1 - \mu_2)} \|g_1^N - g_2^N\|^2.$$

By incorporating this into the telescoped sum and once again taking the minimum over all (sub)gradients residual norms, we obtain the second rate from Corollary 4.2.1. Furthermore, note that a necessary condition for the tightness of these sublinear rates is that the subgradients residual norms $\|g_1^k - g_2^k\|$ must be equal for any $k = 0, \ldots, N$. □

Proofs for more iterations for residual gradient

For simplicity of notation, in this subsection we consider $\eta = \frac{L_2}{\mu_1}$ and $\rho = \frac{\mu_2}{\mu_1}$. Condition $\eta + \rho \leq 0$ is equivalent to $L_2 + \mu_2 \leq 0$.

Recall that Theorem 4.2.2 shows the rate in the domain of regime p_4, while Theorem 4.2.3 gives the rate in the subdomain of regime p_5 with $L_2 + \mu_2 \leq 0$.

To prove these results, we establish additional one-step descent lemmas. Lemma 4.5.1 and Lemma 4.5.2 generalize Lemma 3.5.4 and Lemma 3.5.3, respectively, which were established for GD in Chapter 3 and used to derive the linear regimes in that chapter. Compared with their GD counterparts, they introduce two additional mixed-square terms (see the third line of the inequalities below).

Lemma 4.5.1 (DCA one-step descent for residual gradient for linear regime p_4)**.** *Let $f_1 \in \mathcal{F}_{\mu_1,L_1}$ and $f_2 \in \mathcal{F}_{\mu_2,L_2}$, with $L_2 \in [0,\mu_1]$ and $\mu_1+\mu_2>0$. Then after one iteration of DCA starting from x^k it holds that*

$$\Delta F(x^k) \geq -E_k(\eta)\frac{\|G^k\|^2}{2\mu_1} + E_{k+1}(\eta)\frac{\|G^{k+1}\|^2}{2\mu_1} +$$

$$+\frac{-1+(\eta+\rho)E_{k+1}(\eta)}{\eta-\rho}\frac{\|G^{k+1}-L_2\Delta x^k\|^2}{2\mu_1} +$$

$$+\frac{L_1+\mu_1}{L_1-\mu_1}E_k(\eta)\frac{\|G^k-\mu_1\Delta x^k\|^2}{2\mu_1} + \frac{\|G^k-\mu_1\Delta x^k\|^2}{2(L_1-\mu_1)}.$$

Moreover, if $\mu_2 \in \big[\frac{-L_2\mu_1}{L_2+\mu_1}, L_2\big)$, then all multipliers of the mixed squares are nonnegative.

Proof. The inequality is obtained by summing (4.15) with (4.17) multiplied by $E_{k+1}(\eta)$, and (4.16) multiplied by $-1+\eta E_{k+1}(\eta)$. Moreover, under $\mu_2 \geq \frac{-L_2\mu_1}{L_2+\mu_1}$, we have

$$-1+\frac{L_2+\mu_2}{\mu_1}E_{k+1}(\eta) \geq -1+\frac{L_2+\mu_2}{\mu_1}E_1(\eta) = \frac{L_2\mu_2+\mu_1L_2+\mu_1\mu_2}{L_2^2} \geq 0.$$

The first inequality follows from the monotonicity of $E_k(\eta)$, while the second follows from the lower bound on μ_2. Note that this expression is exactly $L_2\mu_1(\mu_2 B)$. □

Proof of Theorem 4.2.2. The curvature assumptions imply that all mixed terms have nonnegative coefficients and can therefore be dropped, yielding

$$\Delta F(x^k) \geq -E_k(\eta)\frac{\|G^k\|^2}{2\mu_1} + E_{k+1}(\eta)\frac{\|G^{k+1}\|^2}{2\mu_1}.$$

The conclusion follows by telescoping the inequality for $k=0,\dots,N-1$. □

Lemma 4.5.2 (DCA one-step descent for residual gradient for a linear subdomain of regime p_5)**.** *Let $f_1 \in \mathcal{F}_{\mu_1,L_1}$ and $f_2 \in \mathcal{F}_{\mu_2,L_2}$, with $\mu_2 \leq 0$*

and $\mu_1 + \mu_2 > 0$*. Then after one iteration of DCA starting from* x^k *it holds that*

$$\Delta F(x^k) \geq -E_k(\rho)\frac{\|G^k\|^2}{2\mu_1} + E_{k+1}(\rho)\frac{\|G^{k+1}\|^2}{2\mu_1} + \\ + \frac{1-(\eta+\rho)E_{k+1}(\rho)}{\eta-\rho}\frac{\|G^{k+1}-\mu_2\Delta x^k\|^2}{2\mu_1} + \\ + \frac{L_1+\mu_1}{L_1-\mu_1}E_k(\rho)\frac{\|G^k-\mu_1\Delta x^k\|^2}{2\mu_1} + \frac{\|G^k-\mu_1\Delta x^k\|^2}{2(L_1-\mu_1)}.$$

Moreover, if $L_2 + \mu_2 \leq 0$*, then all multipliers of the mixed squares are nonnegative.*

Proof. The inequality is obtained by summing (4.15) with (4.17) multiplied by $E_{k+1}(\rho)$, and (4.16) multiplied by $(-\rho)E_{k+1}(\rho)$. Note that $E_{k+1}(\rho)$ is positive and monotonically increasing with index k as $\mu_1 + \mu_2 > 0$. Then, due to $L_2 + \mu_2 \leq 0$ (equivalent to $\eta + \rho \leq 0$), implying the conclusion. □

Proof of Theorem 4.2.3. The curvature assumptions imply that all mixed squares in Lemma 4.5.2 have nonnegative coefficients and therefore can be dropped, yielding

$$\Delta F(x^k) \geq -E_k(\rho)\frac{\|G^k\|^2}{2\mu_1} + E_{k+1}(\rho)\frac{\|G^{k+1}\|^2}{2\mu_1}.$$

Telescoping the inequality for $k = 0, 1, \ldots, N-1$ leads to the conclusion. □

These proofs closely follow the corresponding GD arguments in Chapter 3, see Section 3.5.2. They require the same multipliers to combine the inequalities, with the additional contribution of f_1 captured in the distance-1 proofs by (4.17).

4.5.2 Proofs of bounds in gradient mapping

We distinguish several types of supporting lemmas for gradient mapping analysis:

- **Fundamental inequalities**: Lemma 4.5.3 establishes the three key distance-1 interpolation inequalities (B^k, $C^k_{f_1}$, $C^k_{f_2}$) that form the basis for all gradient mapping proofs (Section 4.5.2).

- **Individual regime proofs**: Lemmas 4.5.4 to 4.5.9 establish the descent inequalities for regimes r_1 through r_6, respectively, through appropriate weighted combinations of the fundamental inequalities (Section 4.5.2).
- **Linear regime descent lemmas**: Lemma 4.5.10 provides the descent inequality for linear convergence in regime r_4, while Lemma 4.5.11 establishes the corresponding result for the subdomain of regime r_5 with $L_2 + \mu_2 \leq 0$ (Section 4.5.2).

Proof of sublinear rate (Theorem 4.3.1)

Lemma 4.5.3 (Inequalities for proving bounds in gradient mapping)**.** *Let $f_1 \in \mathcal{F}_{\mu_1,L_1}$ and $f_2 \in \mathcal{F}_{\mu_2,L_2}$, with $L_2 > \mu_2$ and $L_1 > \mu_1 \geq 0$. Consider one DCA iteration connecting x^k and x^{k+1}. Let $g_1^j \in \partial f_1(x^j)$, $g_2^j \in \partial f_2(x^j)$ and $G^j = g_1^j - g_2^j$, where $j = \{k, k+1\}$. The following inequalities hold:*

$$\Delta F(x^k) \geq \frac{\mu_1 + \mu_2}{2}\|\Delta x^k\|^2 + \frac{\|G^{k+1} - \mu_2 \Delta x^k\|^2}{2(L_2 - \mu_2)} + \frac{\|G^k - \mu_1 \Delta x^k\|^2}{2(L_1 - \mu_1)}; \qquad (B^k)$$

$$\langle G^k, \Delta x^k \rangle \geq \mu_1 \|\Delta x^k\|^2 + \frac{\|G^k - \mu_1 \Delta x^k\|^2}{L_1 - \mu_1}; \qquad (C^k_{f_1})$$

$$\langle G^{k+1}, \Delta x^k \rangle \geq \mu_2 \|\Delta x^k\|^2 + \frac{\|G^{k+1} - \mu_2 \Delta x^k\|^2}{L_2 - \mu_2}. \qquad (C^k_{f_2})$$

Proof. We exploit the property $g_1^{k+1} = g_2^k$, implying $G^k = g_1^k - g_2^k = g_1^k - g_1^{k+1}$ and $G^{k+1} = g_1^{k+1} - g_2^{k+1} = g_2^k - g_2^{k+1}$. By writing inequality (4.12) for: f_1 with the iterates (x^k, x^{k+1}) and (x^{k+1}, x^k) we get (4.22) and (4.23), respectively; f_2 with the iterates (x^k, x^{k+1}) and (x^{k+1}, x^k) we obtain (4.24) and (4.25), respectively:

$$f_1^k - f_1^{k+1} - \langle g_1^{k+1}, \Delta x^k \rangle \geq \frac{\mu_1}{2}\|\Delta x^k\|^2 + \frac{\|G^k - \mu_1 \Delta x^k\|^2}{2(L_1 - \mu_1)}; \qquad (4.22)$$

$$f_1^{k+1} - f_1^k + \langle g_1^k, \Delta x^k \rangle \geq \frac{\mu_1}{2}\|\Delta x^k\|^2 + \frac{\|G^k - \mu_1 \Delta x^k\|^2}{2(L_1 - \mu_1)}; \qquad (4.23)$$

$$f_2^k - f_2^{k+1} - \langle g_2^{k+1}, \Delta x^k \rangle \geq \frac{\mu_2}{2}\|\Delta x^k\|^2 + \frac{\|G^{k+1} - \mu_2 \Delta x^k\|^2}{2(L_2 - \mu_2)}; \qquad (4.24)$$

$$f_2^{k+1} - f_2^k + \langle g_2^k, \Delta x^k \rangle \geq \frac{\mu_2}{2}\|\Delta x^k\|^2 + \frac{\|G^{k+1} - \mu_2 \Delta x^k\|^2}{2(L_2 - \mu_2)}. \qquad (4.25)$$

Then summing (4.22) and (4.25) gives (B^k); inequality ($C^k_{f_1}$) results by summing (4.22) and (4.23); inequality ($C^k_{f_2}$) is obtained by summing (4.24) and (4.25). ☐

By neglecting the mixed squares in inequality (B^k) it results an alternative proof of the sufficient decrease property from Proposition 4.1.2: $F(x^k) - F(x^{k+1}) \geq (\mu_1 + \mu_2)\frac{1}{2}\|x^k - x^{k+1}\|^2$.

Remark 4.5.2 (Proofs structure for gradient mapping)**.** *The proofs are based on multiplying the inequalities from Lemma 4.5.3 with nonnegative weights. For regimes $r_{1,3,4,5}$ we sum:*

$$B^k + \beta_1 C^{k+1}_{f_1} + \beta_2 C^k_{f_2}, \tag{4.26}$$

while for the proofs of $r_{2,6}$ we sum:

$$B^{k+1} + \beta_1 C^k_{f_1} + \beta_2 C^{k+1}_{f_2},$$

where $B^{\{k,k+1\}}$, $C^{\{k,k+1\}}_{f_1}$ and $C^{\{k,k+1\}}_{f_2}$ are instances of inequalities (B^k), ($C^k_{f_1}$) and ($C^k_{f_2}$), respectively. The weights β_1 and β_2 are functions of the curvature parameters.

The demonstrations rewrite the inequalities from Remark 4.5.2 by building specific mixed squared norms of G and Δx at specific indices k, using Property 4.5.1. Since the derivations are very similar to the ones for residual gradient (given in Section 4.5.1), we only give the final inequalities containing all the squares.

Property 4.5.1. *The following identities hold, where $i, j \in \{1, 2\}$ and G and Δx are specialized when completing the squares within the proofs:*

$$-\langle G, \Delta x\rangle = \frac{1}{2\mu_j}\left[\|G - \mu_j \Delta x\|^2 - \|G\|^2 - \|\mu_j \Delta x\|^2\right]$$

and

$$\|G - \mu_i \Delta x\|^2 = \frac{\mu_i}{\mu_j}\|G - \mu_j \Delta x\|^2 - \frac{\mu_i}{\mu_j}(1 - \frac{\mu_i}{\mu_j})\|\mu_j \Delta x\|^2 + (1 - \frac{\mu_i}{\mu_j})\|G\|^2.$$

Lemma 4.5.4 (Proof of regime r_1)**.** *Let $f_1 \in \mathcal{F}_{\mu_1,L_1}$ and $f_2 \in \mathcal{F}_{\mu_2,L_2}$, with $L_2 > \mu_1 \geq \mu_2 \geq 0$. Then after two DCA iterations connecting x^k, x^{k+1} and x^{k+2} it holds*

$$\Delta F(x^k) \geq \Big(\mu_1 + \mu_2 \frac{L_2 - \mu_1}{L_2 - \mu_2}\Big)\frac{1}{2}\|\Delta x^k\|^2 + \mu_1 \frac{\mu_1 - \mu_2}{L_2 - \mu_2}\frac{1}{2}\|\Delta x^{k+1}\|^2,$$

with equality only if $G^k = G^{k+1} = \mu_1 \Delta x^k = \mu_1 \Delta x^{k+1}$.

Proof. Let $\beta_2 = 0$ and $\beta_1 := \frac{\mu_1-\mu_2}{L_2-\mu_2}$, which is nonnegative since $\mu_1 \geq \mu_2$. After replacing in (4.26), completing squares and performing simplifications we get:

$$\Delta F(x^k) \geq \Big(\mu_1 + \mu_2 \frac{L_2-\mu_1}{L_2-\mu_2}\Big)\frac{1}{2}\|\Delta x^k\|^2 + \mu_1 \frac{\mu_1-\mu_2}{L_2-\mu_2}\frac{1}{2}\|\Delta x^{k+1}\|^2 +$$
$$\frac{\|G^k - \mu_1 \Delta x^k\|^2}{2(L_1-\mu_1)} + \frac{\mu_2\|G^{k+1} - \mu_1\Delta x^k\|^2}{2\mu_1(L_2-\mu_2)} +$$
$$\frac{(L_1+\mu_1)(\mu_1-\mu_2)\|G^{k+1}-\mu_1\Delta x^{k+1}\|^2}{2\mu_1(L_1-\mu_1)(L_2-\mu_2)}.$$

Since $\mu_1 \geq \mu_2 \geq 0$, all coefficients of the squares are positive and the conclusion follows by dropping the mixed terms, with equality only if $G^k = \mu_1 \Delta x^k$ and $G^{k+1} = \mu_1 \Delta x^k = \mu_1 \Delta x^{k+1}$, which in particular also implies $\Delta x^k = \Delta x^{k+1}$. □

Lemma 4.5.5 (Proof of regime r_2)**.** *Let $f_1 \in \mathcal{F}_{\mu_1,L_1}$ and $f_2 \in \mathcal{F}_{\mu_2,L_2}$, with $L_1 > \mu_2 \geq \mu_1 \geq 0$. Then after two DCA iterations connecting x^k, x^{k+1} and x^{k+2} it holds*

$$\Delta F(x^{k+1}) \geq \mu_2 \frac{\mu_2-\mu_1}{L_1-\mu_1}\frac{1}{2}\|\Delta x^k\|^2 + \Big(\mu_2 + \mu_1\frac{L_1-\mu_2}{L_1-\mu_1}\Big)\frac{1}{2}\|\Delta x^{k+1}\|^2,$$

with equality only if $G^{k+1} = G^{k+2} = \mu_2\Delta x^k = \mu_2 \Delta x^{k+1}$.

Proof. Let $\beta_1 = 0$ and $\beta_2 := \frac{\mu_2-\mu_1}{L_1-\mu_1}$, nonnegative since $\mu_2 \geq \mu_1$, be the specific multipliers from Remark 4.5.2. Completing the mixed squares we obtain:

$$\Delta F(x^{k+1}) \geq \Big(\mu_2 + \mu_1\frac{L_1-\mu_2}{L_1-\mu_1}\Big)\frac{1}{2}\|\Delta x^{k+1}\|^2 + \mu_2\frac{\mu_2-\mu_1}{2(L_1-\mu_1)}\|\Delta x^k\|^2 +$$
$$\frac{\|G^{k+2}-\mu_2\Delta x^{k+1}\|^2}{2(L_2-\mu_2)} + \frac{\mu_1\|G^{k+1}-\mu_2\Delta x^{k+1}\|^2}{2\mu_2(L_1-\mu_1)} +$$
$$\frac{(L_2+\mu_2)(\mu_2-\mu_1)\|G^{k+1}-\mu_2\Delta x^k\|^2}{2\mu_2(L_2-\mu_2)(L_1-\mu_1)}.$$

Since $\mu_2 \geq \mu_1 \geq 0$, all coefficients of the squares are positive and the conclusion follows by dropping the mixed terms, with equality only if $G^{k+2} = \mu_2\Delta x^{k+1}$ and $G^{k+1} = \mu_2 \Delta x^k = \mu_2\Delta x^{k+1}$, also implying $\Delta x^k = \Delta x^{k+1}$. □

A quantity delimiting the regimes is T_1, as defined in Property 4.5.2. It is exactly zero when $B = 0$ (where $B = \mu_1^{-1} + \mu_2^{-2} + L_2^{-1}$ also delimits the regimes in the residual gradient case).

Property 4.5.2 (T_1). *When $\mu_2 < 0$ and $\mu_1 > 0$, we define*

$$T_1 := E_N\left(\frac{L_2}{\mu_1}\right) - E_1\left(\frac{\mu_2}{\mu_1}\right) = \mu_1\Big(\frac{L_2+\mu_1}{L_2^2} - \frac{\mu_1+\mu_2}{\mu_2^2}\Big).$$

One can check that $T_1 \geq 0$ if $\mu_2 \in (-\mu_1, \frac{-L_2\mu_1}{L_2+\mu_1}]$ and $T_1 \leq 0$ if $\mu_2 \in [\frac{-L_2\mu_1}{L_2+\mu_1}, L_2\big)$.

Lemma 4.5.6 (Proof of regime r_3). *Let $f_1 \in \mathcal{F}_{\mu_1,L_1}$ and $f_2 \in \mathcal{F}_{\mu_2,L_2}$, with $L_2 > \mu_1 > 0$, $\mu_2 \in [\frac{-L_2\mu_1}{L_2+\mu_1}, 0\big)$ and $\mu_1 + \mu_2 > 0$. Consider two DCA iterations connecting x^k, x^{k+1} and x^{k+2}. Then*

$$\Delta F(x^k) \geq \Big(\mu_1 + \frac{L_2\mu_2}{L_2+\mu_2}\Big)\frac{1}{2}\|\Delta x^k\|^2 + \frac{\mu_1^2}{2(L_2+\mu_2)}\|\Delta x^{k+1}\|^2,$$

with equality only if $G^k = \mu_1\Delta x^k$ and $G^{k+1} = \mu_1\Delta x^{k+1}$.

Proof. Let $\beta_1 := \frac{\mu_1}{L_2+\mu_2}$ and $\beta_2 := \frac{-\mu_2}{L_2+\mu_2}$, which are nonnegative since $\mu_1 > -\mu_2 > 0$ and $L_2 + \mu_2 \geq \frac{-\mu_2 L_2}{\mu_1} > 0$. After replacing in (4.26), completing squares and performing simplifications we obtain:

$$\Delta F(x^k) \geq \Big(\mu_1 + \frac{L_2\mu_2}{L_2+\mu_2}\Big)\frac{1}{2}\|\Delta x^k\|^2 + \frac{\mu_1^2}{2(L_2+\mu_2)}\|\Delta x^{k+1}\|^2 + \frac{\|G^k - \mu_1\Delta x^k\|^2}{2(L_1-\mu_1)} + \frac{(L_1+\mu_1)\|G^{k+1} - \mu_1\Delta x^{k+1}\|^2}{2(L_1-\mu_1)(L_2+\mu_2)}.$$

Since $L_2 > -\mu_2$, the coefficients of the mixed terms are positive and the conclusion follows by dropping them, also implying the necessary conditions for equality. The coefficient of $\|\Delta x^k\|^2$ is positive if $T_1 \leq 0$ as it can be factorized as $\frac{L_2^2\mu_2^2}{\mu_1(L_2^2-\mu_2^2)}(-T_1)$. □

Lemma 4.5.7 (Proof of regime r_4). *Let $f_1 \in \mathcal{F}_{\mu_1,L_1}$ and $f_2 \in \mathcal{F}_{\mu_2,L_2}$, with $L_2 \in [0, \mu_1]$, $\mu_2 \in \big[\frac{-L_2\mu_1}{L_2+\mu_1}, L_2\big)$ and $\mu_1 + \mu_2 > 0$. Consider two DCA iterations connecting x^k, x^{k+1} and x^{k+2}. Then it holds*

$$\Delta F(x^k) \geq \frac{\mu_1^2(L_2+\mu_1)}{L_2^2}\frac{1}{2}\|\Delta x^{k+1}\|^2,$$

with equality only if $G^k = \mu_1\Delta x^k$ and $G^{k+1} = \mu_1\Delta x^{k+1} = L_2\Delta x^k$.

Proof. Let $\beta_1 := \frac{\mu_1(\mu_1+L_2)}{L_2^2}$ and $\beta_2 := \frac{\mu_1}{L_2}$, both positive. After replacing in (4.26), completing the mixed squares and performing simplifications we get:

$$\Delta F(x^k) \geq \frac{\mu_1^2(L_2+\mu_1)}{L_2^2}\frac{1}{2}\|\Delta x^{k+1}\|^2 + \frac{\|G^k - \mu_1\Delta x^k\|^2}{2(L_1-\mu_1)} +$$

$$\frac{(L_1+\mu_1)(L_2+\mu_1)\|G^{k+1}-\mu_1\Delta x^{k+1}\|^2}{L_2^2(L_1-\mu_1)}+\frac{\mu_2^2\,(-T_1)\|G^{k+1}-L_2\Delta x^k\|^2}{\mu_1(L_2-\mu_2)^2}.$$

Since $T_1 \le 0$ (Property 4.5.2), the conclusion follows by dropping the mixed terms. □

Lemma 4.5.8 (Proof of regime r_5). *Let $f_1 \in \mathcal{F}_{\mu_1,L_1}$ and $f_2 \in \mathcal{F}_{\mu_2,L_2}$, with $\mu_1 > -\mu_2 > 0$ such that $\mu_2 \in \big(-\mu_1, \frac{-L_2\mu_1}{L_2+\mu_1}\big]$. Consider two DCA iterations connecting x^k, x^{k+1} and x^{k+2}. Then it holds*

$$\Delta F(x^k) \ge \frac{\mu_1^2(\mu_1+\mu_2)}{\mu_2^2}\frac{1}{2}\|\Delta x^{k+1}\|^2,$$

with equality only if $G^k = \mu_1\Delta x^k$ and $G^{k+1} = \mu_1\Delta x^{k+1} = \mu_2\Delta x^k$.

Proof. Let $\beta_1 := \frac{\mu_1(\mu_1+\mu_2)}{\mu_2^2}$ and $\beta_2 := \frac{\mu_1+\mu_2}{-\mu_2}$, both positive. After replacing in (4.26), completing the mixed squares and performing simplifications we get:

$$\Delta F(x^k) \ge \frac{\mu_1^2(\mu_1+\mu_2)}{2\mu_2^2}\|\Delta x^{k+1}\|^2+\frac{\|G^k-\mu_1\Delta x^k\|^2}{2(L_1-\mu_1)}+$$
$$\frac{(L_1+\mu_1)(\mu_1+\mu_2)\|G^{k+1}-\mu_1\Delta x^{k+1}\|^2}{\mu_2^2(L_1-\mu_1)}+\frac{L_2^2\,T_1\|G^{k+1}-\mu_2\Delta x^k\|^2}{\mu_1(L_2-\mu_2)^2}.$$

Since $T_1 \ge 0$ (Property 4.5.2), the conclusion follows by dropping all mixed squares terms. □

Lemma 4.5.9 (Proof of regime r_6). *Let $f_1 \in \mathcal{F}_{\mu_1,L_1}$ and $f_2 \in \mathcal{F}_{\mu_2,L_2}$, with $\mu_2 \ge \mu_1 \ge 0$ and $L_1 \in (0,\mu_2]$. Then after two DCA iterations connecting x^k, x^{k+1} and x^{k+2} we have*

$$\Delta F(x^{k+1}) \ge \frac{\mu_2^2(L_1+\mu_2)}{L_1^2}\frac{1}{2}\|\Delta x^k\|^2,$$

with equality only if $G^{k+2} = \mu_2\Delta x^{k+1}$ and $G^{k+1} = \mu_2\Delta x^k = L_1\Delta x^{k+1}$.

Proof. Let $\beta_2 := \frac{\mu_2(\mu_2+L_1)}{L_1^2}$ and $\beta_1 := \frac{\mu_2}{L_1}$, both positive. After replacing in (4.26), completing the mixed squares and performing simplifications we get:

$$\Delta F(x^{k+1}) \ge \frac{\mu_2^2(L_1+\mu_2)}{L_1^2}\frac{1}{2}\|\Delta x^k\|^2+\frac{\|G^{k+2}-\mu_2\Delta x^{k+1}\|^2}{2(L_2-\mu_2)}+$$
$$\frac{(L_1+\mu_2)(L_2+\mu_2)\|G^{k+1}-\mu_2\Delta x^k\|^2}{L_1^2(L_2-\mu_2)}+$$

$$\frac{\mu_1\mu_2 + L_1(\mu_1+\mu_2)\|G^{k+1} - L_1\Delta x^{k+1}\|^2}{L_1^2(L_1-\mu_1)}.$$

The conclusion follows by dropping all mixed squares terms. □

Proof of Theorem 4.3.1. For regimes r_1, r_3, r_4 and r_5, by telescoping inequalities from Lemma 4.5.4, Lemma 4.5.6, Lemma 4.5.7 and Lemma 4.5.8, respectively, for $k = \{0, \dots, N-1\}$ and taking the minimum over iterations differences norms we get:

$$F(x^0) - F(x^N) \geq r_i(\mu_1, L_1, \mu_2, L_2)N\,\frac{1}{2}\min_{0\leq k\leq N}\{\|x^k - x^{k+1}\|^2\},$$

where r_i, with $i = \{1, 3, 4, 5\}$, denotes the sum of the two coefficients multiplying $\frac{1}{2}\|x^k - x^{k+1}\|^2$ and $\frac{1}{2}\|x^{k+1} - x^{k+2}\|^2$ in each inequality. Writing (5.3) for the iterates x^N and x^{N+1}, we have $F(x^N) - F(x^{N+1}) \geq \frac{\mu_1+\mu_2}{2}\|x^N - x^{N+1}\|^2$. Then the rate results by summing it to the previous inequality and taking the minimum in the right-hand side. For regimes r_2 and r_6, by telescoping inequalities from Lemmas 4.5.5 and 4.5.9, respectively, for $k = \{1, \dots, N\}$ and taking the minimum over iterations:

$$F(x^1) - F(x^{N+1}) \geq r_i(\mu_1, L_1, \mu_2, L_2)N\,\frac{1}{2}\min_{0\leq k\leq N}\{\|x^k - x^{k+1}\|^2\},$$

with $i = \{2, 6\}$. Writing (5.3) for the iterates x^0 and x^1, summing it up and taking the minimum in the right-hand side leads to the rate. □

We remark the symmetry between regimes $r_1 \iff r_2$ and $r_4 \iff r_6$, as the latter ones reflect the behaviours of the former regimes, as the iterations would be applied in reverse order on the function $-F = f_2 - f_1$. In other words, Lemmas 4.5.5 and 4.5.9 can be obtained from Lemmas 4.5.4 and 4.5.7, respectively, by interchanging the curvatures $\mu_1 \iff \mu_2$ and $L_1 \iff L_2$ and the iterations order, i.e., from $k+2$ to k. This symmetry between the proofs was also remarked for residual gradient metric in Section 4.5.1.

Moreover, the proofs of sublinear rates for regimes $r_{1,3,4,5}$ are obtained by *appending* the sum of their corresponding lemmas after N iterations, whereas the proofs for regimes $r_{2,6}$ result by *prepending* the telescoped sum.

Proofs for more iterations for gradient mapping

For simplicity of notation, in this subsection we consider $\eta = \frac{L_2}{\mu_1}$ and $\rho = \frac{\mu_2}{\mu_1}$. Condition $\eta + \rho \leq 0$ is equivalent to $L_2 + \mu_2 \leq 0$. The proofs are very similar to what we have for residual gradient in Section 4.5.1 (see Lemmas 4.5.1 and 4.5.2) and are alternative generalizations of Lemma 3.5.4 and Lemma 3.5.3.

Lemma 4.5.10 (DCA descent lemma for gradient mapping for linear regime r_4)**.** *Let $f_1 \in \mathcal{F}_{\mu_1,L_1}$ and $f_2 \in \mathcal{F}_{\mu_2,L_2}$, with $L_2 \in [0,\mu_1]$ and $\mu_1+\mu_2>0$. Then after two DCA iterations starting from x^k it holds that*

$$\Delta F(x^k) \geq -E_k(\eta)\tfrac{\mu_1}{2}\|\Delta x^k\|^2 + E_{k+1}(\eta)\tfrac{\mu_1}{2}\|\Delta x^{k+1}\|^2 + \\ + \frac{-1+(\eta+\rho)E_{k+1}(\eta)}{\eta-\rho}\frac{\|G^{k+1}-L_2\Delta x^k\|^2}{2\mu_1} + \\ + \frac{\|G^k-\mu_1\Delta x^k\|^2}{2(L_1-\mu_1)} + \frac{L_1+\mu_1}{L_1-\mu_1}E_{k+1}(\eta)\frac{\|G^{k+1}-\mu_1\Delta x^{k+1}\|^2}{2\mu_1}.$$

Moreover, if $\mu_2 \in \left[\frac{-L_2\mu_1}{L_2+\mu_1}, L_2\right)$, then all multipliers of the mixed squares are nonnegative.

Proof. The inequality is obtained by taking $\beta_1 = E_{k+1}(\eta)$ and $\beta_2 = -1 + \eta E_{k+1}(\eta)$ in (4.26) and completing the squares. Moreover, under $\mu_2 \geq \frac{-L_2\mu_1}{L_2+\mu_1}$ and since $E_{k+1}(\eta)$ is increasing with k, we have that

$$-1+\frac{L_2+\mu_2}{\mu_1}E_{k+1}(\eta) \geq -1+\frac{L_2+\mu_2}{\mu_1}E_1(\eta) = \frac{\mu_2^2}{\mu_1(L_2-\mu_2)}(-T_1) \geq 0,$$

thus the mixed terms have nonnegative coefficients. □

Proof of Theorem 4.3.2. The curvature assumptions imply that all mixed squares in Lemma 4.5.10 have nonnegative coefficients and therefore can be dropped, yielding

$$\Delta F(x^k) \geq -E_k(\eta)\tfrac{\mu_1}{2}\|\Delta x^k\|^2 + E_{k+1}(\eta)\tfrac{\mu_1}{2}\|\Delta x^{k+1}\|^2.$$

Telescoping for $k=0,\dots,N-1$ and adding (B^k) with $k=N$ gives the conclusion. □

Lemma 4.5.11 (DCA descent lemma for gradient mapping for a linear subdomain of regime r_5)**.** *Let $f_1 \in \mathcal{F}_{\mu_1,L_1}$ and $f_2 \in \mathcal{F}_{\mu_2,L_2}$, with $\mu_2 \leq 0$ and $\mu_1+\mu_2>0$. Then after two DCA iterations starting from x^k it holds that*

$$\Delta F(x^k) \geq -E_k(\rho)\frac{\mu_1}{2}\|\Delta x^k\|^2 + E_{k+1}(\rho)\frac{\mu_1}{2}\|\Delta x^{k+1}\|^2 + \\ + \frac{1-(\eta+\rho)E_{k+1}(\rho)}{\eta-\rho}\frac{\|G^{k+1}-\mu_2\Delta x^k\|^2}{2\mu_1} + \\ + \frac{L_1+\mu_1}{L_1-\mu_1}E_{k+1}(\rho)\frac{\|G^{k+1}-\mu_1\Delta x^{k+1}\|^2}{2\mu_1} + \frac{\|G^k-\mu_1\Delta x^k\|^2}{2(L_1-\mu_1)}.$$

Moreover, if $L_2 + \mu_2 \leq 0$, then all multipliers of the mixed squares are nonnegative.

Proof. The inequality is obtained by taking $\beta_1 = E_{k+1}(\rho)$ and $\beta_2 = -\rho E_{k+1}(\rho)$ in (4.26) and completing the squares. Moreover, since $E_{k+1}(\rho)$ is positive and monotonically increasing with index k as $\mu_1 + \mu_2 > 0$, and $L_2 + \mu_2 \leq 0$, the mixed squares have nonnegative coefficients. □

Proof of Theorem 4.3.3. The curvature assumptions imply that all mixed squares in Lemma 4.5.11 can be dropped because they have nonnegative coefficients. This yields the inequality

$$\Delta F(x^k) \geq -E_k(\rho)\frac{\mu_1}{2}\|\Delta x^k\|^2 + E_{k+1}(\rho)\frac{\mu_1}{2}\|\Delta x^{k+1}\|^2.$$

Telescoping for $k = 0, \ldots, N-1$ and adding (B^k) with $k = N$ gives the conclusion. □

Lemmas 4.5.10 and 4.5.11 closely parallel Lemmas 4.5.1 and 4.5.2, respectively, with the following modifications: they track the iterate displacement $\mu_1 \Delta x^{k\,/\,k+1}$ rather than $G^{k\,/\,k+1}$, and they involve a mixed term at iteration $k+1$ (namely, $E_{k+1}(\cdot)\,\|G^{k+1} - \mu_1 \Delta x^{k+1}\|^2$) instead of at iteration k. In turn, Lemmas 4.5.10 and 4.5.11 also generalize Lemmas 3.5.3 and 3.5.4, respectively, which were established for GD in Chapter 3.

Remark 4.5.3 (Recovering GD descent lemmas from DCA analysis)**.** *Anticipating the PGD–DCA equivalence results in Section 4.6, the corresponding GD lemmas are recovered by enforcing the GD relation $G^k = \mu_1 \Delta x^k$ for all k (which eliminates the last two mixed terms) and by making the following curvature substitutions (on the r.h.s. we use the GD notation from Chapter 3):*

$$\mu_1 = \gamma^{-1}, \qquad \mu_2 = \gamma^{-1} - L, \qquad \eta = 1 - \gamma\mu, \qquad \rho = 1 - \gamma L.$$

4.5.3 Proof of bounds for both nonsmooth

Proof of Theorem 4.4.1. The subgradient inequality of f_2 reads:

$$f_2(x^+) \geq f_2(x) + \langle g_2, x^+ - x\rangle + \tfrac{1}{2}\mu_2\|x - x^+\|^2. \tag{4.27}$$

Case $\mu_2 \geq 0$. By adding and subtracting $f_1(x) - f_1(x^+)$ in both sides of (4.27) we obtain exactly (4.9):

$$F(x) - F(x^+) \geq f_1(x) - f_1(x^+) + \langle g_2, x^+ - x\rangle + \tfrac{1}{2}\mu_2\|x - x^+\|^2.$$

Case $\mu_2 < 0$**.** We have $\mu_1 > -\mu_2 > 0$. From the strong convexity of f_1, with $g_1^+ \in \partial f_1(x)$:

$$f_1(x) \geq f_1(x^+) - \langle g_1^+, x^+ - x\rangle + \tfrac{1}{2}\mu_1\|x - x^+\|^2.$$

By multiplying it with $-\frac{\mu_2}{\mu_1} > 0$ and summing with (4.27) we get, using $g_1^+ = g_2$:

$$-\frac{\mu_2}{\mu_1}[f_1(x) - f_1(x^+)] - f_2(x) + f_2(x^+) \geq \big(1 + \frac{\mu_1}{\mu_2}\big)\langle g_2, x^+ - x\rangle.$$

By adding $(1 + \frac{\mu_2}{\mu_1})(f_1(x) - f_1(x^+))$ in both sides we obtain exactly (4.10):

$$F(x) - F(x^+) \geq \big(1 + \frac{\mu_1}{\mu_2}\big)\left[f_1(x) - f_1(x^+) + \langle g_2, x^+ - x\rangle\right].$$

□

4.6 Convergence rates for proximal gradient descent

In this section, we show that the convergence rates of the proximal gradient descent (PGD) method can be directly derived from its iterate-wise equivalence with the DCA, a connection that is often underemphasized in the literature. Although DCA is typically applied to nonconvex-nonconcave objective functions, it also provides a valuable framework for establishing performance bounds in problems involving (strongly) convex objectives. Standard references for the fundamental theory of proximal methods include [81, 97, 78, 13, 30, 93, 15].

Assumption 4.6.1 (PGD splitting setup)**.** *Consider the composite objective function* $F = \varphi + h$*, where* $\varphi \in \mathcal{F}_{\mu_\varphi, L_\varphi}$ *is smooth, with* $L_\varphi > 0$ *and* $\mu_\varphi \in (-\infty, L_\varphi)$*, and* $h\colon \mathbb{R}^d \to (-\infty, \infty]$ *is proper, closed and convex,* $h \in \mathcal{F}_{\mu_h, L_h}$*, such that* $0 \leq \mu_h \leq L_h$*.*

The PGD iteration with stepsize $\gamma > 0$, starting from x^k, is given by:

$$\begin{aligned} x^{k+1} &:= \operatorname{prox}_{\gamma h}\left[x^k - \gamma\nabla\varphi(x^k)\right] \\ &= \operatorname*{arg\,min}_{w \in \mathbb{R}^d} \left\{h(w) + \frac{1}{2\gamma}\left\|w - x^k + \gamma\nabla\varphi(x^k)\right\|^2\right\}. \end{aligned} \tag{PGD}$$

Drori established in [41] a tight bound in the convex case of projected gradient descent, measuring the optimality gap $F(x^N) - F_*$. When φ is (strongly) convex, Taylor et al. [114] establish tight rates for any constant stepsize and curvature choice when employing different performance metrics than ours. In the case of our residual gradient measure, they give the rate for the stepsize choice $\gamma = \frac{1}{L_\varphi}$ (see Table 2 therein), letting as further work the rest of stepsize interval (see Table 1); this work is completed by us. Additionally, they do not consider the gradient mapping measure.

When φ is smooth nonconvex ($\mu_\varphi = -L_\varphi$) and h is convex, Li and Pong [77] show convergence to critical points when using stepsizes $\gamma < \frac{1}{L_\varphi}$. Measuring the residual gradient, considering φ nonconvex, with $\mu_\varphi = -L_\varphi$ and $\mu_h = 0$, $L_h = \infty$, the tight rate for stepsizes shorter than $\frac{1}{L_\varphi}$ is given in [1, Proposition 8.11]. For other cases in the smooth weakly convex setting, we are not aware of previous tight convergence results.

Remark 4.6.1 (On residual gradient vs gradient mapping)**.** *There is a subtle difference between residual gradient and gradient mapping, which leads us to develop two different lines of proofs. With $g_h^k \in \partial h(x^k)$, the residual gradient is measured as $\min_{0\leq k\leq N}\{\|\nabla\varphi(x^k) + g_h^k\|\}^2$, whereas the gradient mapping is measured as $\min_{0\leq k\leq N}\{\|\gamma^{-1}(\nabla\varphi(x^k) + g_h^{k+1})\|^2\}$ (this expression is obtained after replacing the proximal mapping optimality condition). What differs is the point at which the (sub)gradient of the (usually nonsmooth) function h is evaluated, namely either at the current iterate x^k or the next one, x^{k+1}.*

Leveraging the equivalence between PGD and DCA, one can directly get the rates for PGD from the DCA ones for any admissible curvature choice. Further on, we focus on several particular cases, with a complementary cover of the rates given in residual gradient (in Section 4.6.2) and gradient mapping (in Section 4.6.3). Finally, in Section 4.6.4 we provide performance bounds for the proximal point algorithm (PPA).

4.6.1 PGD iteration as a special case of DCA iteration

In Proposition 4.6.1, we show that one iteration of PGD can be viewed as one iteration of DCA when f_2 is smooth. To the best of our knowledge, the first formal proof of this equivalence is due to Thi and Pham Dinh [74, §3.3.4].

Proposition 4.6.1 (PGD is equivalent to DCA [74, §3.3.4])**.** *Starting from $x \in \mathbb{R}^d$, one iteration of* (PGD) *with stepsize $\gamma > 0$ on the composite objective function $F = \varphi + h$, under Assumption 4.6.1, yields the same point x^+ as a DCA iteration applied to the splitting $F = f_1 - f_2$, where $f_1 = h + \frac{1}{2\gamma}\|\cdot\|^2$ and $f_2 = \frac{1}{2\gamma}\|\cdot\|^2 - \varphi$. Furthermore, the curvatures are related as follows: $\mu_1 = \gamma^{-1} + \mu_h$; $L_1 = \gamma^{-1} + L_h$; $\mu_2 = \gamma^{-1} - L_\varphi$; $L_2 = \gamma^{-1} - \mu_\varphi$.*

Proof. The approach is to reformulate a PGD iteration as a DCA one, starting from the proximal step definition:

$$\begin{aligned} x^+ &= \operatorname{prox}_{\gamma h}\left[x-\gamma\nabla\varphi(x)\right] \\ &= \operatorname*{arg\,min}_{w\in\mathbb{R}^d}\left\{h(w)+\frac{1}{2\gamma}\left\|w-x+\gamma\nabla\varphi(x)\right\|^2\right\} \\ &= \operatorname*{arg\,min}_{w\in\mathbb{R}^d}\left\{h(w)+\frac{\left\|w\right\|^2}{2\gamma}+\left\langle\nabla\varphi(x)-\frac{1}{\gamma}x,w\right\rangle\right\} \\ &= \operatorname*{arg\,min}_{w\in\mathbb{R}^d}\left\{\left[h(w)+\frac{\left\|w\right\|^2}{2\gamma}\right]-\left\langle\nabla\left[\frac{\|x\|^2}{2\gamma}-\varphi(x)\right],w\right\rangle\right\}. \end{aligned}$$

Then, from $f_1 = h + \frac{1}{2\gamma}\|\cdot\|^2$ and $f_2 = \frac{1}{2\gamma}\|\cdot\|^2 - \varphi$, we recover the (DCA) iteration for f_2 smooth. □

Table 4.6.1: Equivalence of curvatures between DCA and PGD as shown in Proposition 4.6.1.

	DCA	$\iff$	**PGD**
Objective	$F\in\mathcal{F}_{\mu_1-L_2,L_1-\mu_2}$	$\iff$	$F\in\mathcal{F}_{\mu_\varphi+\mu_h,L_\varphi+L_h}$
Curvatures	μ_1	$\iff$	$\gamma^{-1}+\mu_h$
	L_1	$\iff$	$\gamma^{-1}+L_h$
	μ_2	$\iff$	$\gamma^{-1}-L_\varphi$
	L_2	$\iff$	$\gamma^{-1}-\mu_\varphi$
Convergence condition	$\mu_1+\mu_2>0$	$\iff$	$\gamma<\dfrac{2}{L_\varphi-\mu_h}$
Standard DC	$\mu_2\geq 0$	$\iff$	$\gamma\leq\dfrac{1}{L_\varphi}$

Proposition 4.6.2 (Equivalence of residual gradient criteria in PGD and DCA)**.** *The convergence measure based on the (sub)gradients residual norm is identical when applying PGD to the decomposition $F=\varphi+h$ or DCA to the decomposition $F=f_1-f_2$, where f_1 and f_2 are defined in Proposition 4.6.1.*

Proof. For any $x \in \mathbb{R}^d$, let $g_h \in \partial h(x)$ and define $G \coloneqq \nabla\varphi(x) + g_h \in \partial F(x)$. By Proposition 4.6.1, we have that $g_1 \coloneqq g_h + \frac{x}{\gamma} \in \partial f_1(x)$. Moreover, the smoothness of φ implies the one of f_2. Then the following holds:

$$G = \nabla\varphi(x) + g_h = \left[\frac{x}{\gamma} - \nabla f_2(x)\right] + \left[g_1 - \frac{x}{\gamma}\right] = g_1 - \nabla f_2(x).$$

Consequently, $\|G\|^2 = \|\nabla\varphi(x) + g_h\|^2 = \|g_1 - \nabla f_2(x)\|^2$. Since F is unchanged between the two splittings, the conclusion follows. □

Proposition 4.6.2 shows that for any iteration x it holds $\|\nabla\varphi(x) + g_h\|^2 = \|g_1 - \nabla f_2(x)\|^2$, where $g_h \in \partial h(x)$ and $g_1 \in \partial f_1(x)$. To determine convergence rates for PGD, one can substitute the curvature values from the DCA convergence rate expressions in Section 4.2 with the corresponding curvatures defined in terms of PGD parameters as specified in Proposition 4.6.1 (we emphasize this correspondence in Table 4.6.1). Moreover, the smoothness of φ implies the smoothness of f_2, ensuring that the critical points for which we studied the DCA convergence are also stationary.

Particular cases of PGD. The projected gradient descent (ProjGD) is characterized by setting $h = \delta_C$ (indicator function of a non-empty, closed and convex set C), hence $\mu_h = 0$ and $L_h = \infty$. The proximal point algorithm is obtained by setting φ constant. The gradient descent (GD) is obtained by setting $h = 0$, hence $\mu_h = L_h = 0$. In this case, we further have $x^k - x^{k+1} = -\gamma\nabla\varphi(x^k)$, where $\gamma^{-1} = \mu_1$. Convergence rates on the best gradient norm are proved in [100] for all curvatures of φ, as presented in Chapter 3. These rates are a particular case of Conjecture 4.3.1 and served as a base to conjecture the tight expressions in Section 4.2.3 and Section 4.3.2, since the projection/proximal step is not modifying the behaviour of the worst-case function.

The typical PGD setup involves φ being smooth and h convex and nonsmooth, thus $\mu_h = 0$, $L_h = \infty$. Hence, the DCA-like curvatures are $\mu_1 = \gamma^{-1}$, $L_1 = \infty$, $\mu_2 = \gamma^{-1} - L_\varphi$, $L_2 = \gamma^{-1} - \mu_\varphi$. However, if additional information about h is available, it can be similarly incorporated to derive improved rates.

The running example in Section 1.1 corresponds, for PGD, to the iterative shrinkage–thresholding algorithm (ISTA) [35], obtained by setting $h(x) = \lambda\|x\|_1$ (so that $h \in \mathcal{F}_{0,\infty}$), where $\lambda > 0$ is a regularization parameter.

Furthermore, from the equivalence of curvatures, it follows that large stepsizes $\gamma > \frac{1}{L_\varphi}$ correspond to negative μ_2, indicating that f_2 is weakly convex. Furthermore, the condition $\mu_1 + \mu_2 > 0$ in the DCA setting, ensuring the decrease in the objective after one iteration (see Proposition 4.1.2), translates to the standard upper bound on the stepsize for PGD, which is $\gamma < \frac{2}{L_\varphi}$.

Table 4.6.2: PGD settings in DCA curvatures. $\mu_h = 0$, $L_h = \infty$ imply $\mu_1 = \gamma^{-1} > 0$ and $L_1 = \infty$ (see Proposition 4.6.1) and the corresponding residual gradient regimes.

φ	**Stepsize** γ	$\mu_2 = \gamma^{-1} - L_\varphi$	$L_2 = \gamma^{-1} - \mu_\varphi$	**Regime**
nonconvex $\mu_\varphi < 0$	$\gamma \in \left(0, \frac{1}{L_\varphi}\right)$	$\mu_2 > 0$	$L_2 > \mu_1$	p_1
	$\gamma = \frac{1}{L_\varphi}$	$\mu_2 = 0$		
	$\gamma \in \left(\frac{1}{L_\varphi}, \frac{2}{L_\varphi}\right)$	$\mu_2 < 0$		p_1 or p_5
convex $\mu_\varphi = 0$	$\gamma \in \left(0, \frac{1}{L_\varphi}\right)$	$\mu_2 > 0$	$L_2 = \mu_1$	$p_1 = p_4$
	$\gamma = \frac{1}{L_\varphi}$	$\mu_2 = 0$		
	$\gamma \in \left(\frac{1}{L_\varphi}, \frac{2}{L_\varphi}\right)$	$\mu_2 < 0$		p_5
strongly convex $\mu_\varphi > 0$	$\gamma \in \left(0, \frac{1}{L_\varphi}\right)$	$\mu_2 > 0$	$\lvert\mu_2\rvert < L_2 < \mu_1$	p_4
	$\gamma = \frac{1}{L_\varphi}$	$\mu_2 = 0$		
	$\gamma \in \left(\frac{1}{L_\varphi}, \frac{2}{L_\varphi + \mu_\varphi}\right)$	$\mu_2 < 0$		p_4 or p_5
	$\gamma \in \left[\frac{2}{L_\varphi + \mu_\varphi}, \frac{2}{L_\varphi}\right)$		$L_2 \leq -\mu_2 < \mu_1$	p_5

Moreover, the case of $\mu_2 \geq \mu_1$ corresponds to $L_\varphi \leq 0$, namely concave minimization over a convex set, while the case of $\mu_2 < \mu_1$ corresponds to $L_\varphi > 0$, which translates to convex/nonconvex minimization over a convex set.

4.6.2 Convergence rates in residual gradient

In Table 4.6.2 we summarize the corresponding DCA curvatures for various notable cases in the PGD setting, with the regimes for the residual gradient metric. For the stepsize $\gamma = \frac{1}{L_\varphi}$, commonly used in convergence analysis, we have $\mu_2 = 0$. Moreover, the unusual case where $L_2 \leq 0$, meaning that f_2 is concave, corresponds to large stepsizes $\gamma \in \left[\frac{2}{L_\varphi + \mu_\varphi}, \frac{2}{L_\varphi}\right)$ applied to strongly convex objectives.

Definition 4.6.1 (First stepsize threshold in PGD)**.** *Let* $\overline{\gamma}^1(L_\varphi, \mu_\varphi, \mu_h)$ *denote the root in* $\left(\frac{1}{L_\varphi}, \frac{2}{L_\varphi - \mu_h}\right)$ *of*

$$2 + \gamma(\mu_h - L_\varphi) - \frac{(1 + \gamma\mu_h)^2}{2 - \gamma(\mu_\varphi - \mu_h)} = 0.$$

The equation in Definition 4.6.1 is obtained by substituting the equivalent PGD curvature expressions into the threshold condition $\mu_2 = \frac{-L_2\mu_1}{L_2 + \mu_1}$ (equivalently,

$B = 0$), introduced in the residual gradient analysis. When $\mu_h = 0$, solving this equation yields the threshold $\overline{\gamma L}_1\big(\frac{\mu_\varphi}{L_\varphi}\big)$, which coincides with the first stepsize threshold defined for gradient descent in Section 3.2.5 (see Definition 3.2.2).

Proposition 4.6.3 (PGD one-step descent for the residual gradient with h convex (regimes p_1, p_5))**.** *Let $\varphi \in \mathcal{F}_{\mu_\varphi, L_\varphi}$ and $h \in \mathcal{F}_{0,\infty}$, with φ smooth. Consider one iteration of* (PGD) *with stepsize $\gamma \in \big(0, \frac{2}{L_\varphi}\big)$. Then:*

$$F(x) - F(x^+) \geq \sigma^+(L_\varphi, \mu_\varphi, \gamma)\tfrac{1}{2}\|\nabla\varphi(x^+) + g_h^+\|^2,$$

with $g_h^+ \in \partial h(x^+)$ and $\sigma^+(L_\varphi, \mu_\varphi, \gamma) \geq 0$ defined as:

$$\sigma_i^+(L_\varphi, \mu_\varphi, \gamma) := \begin{cases} \frac{\gamma(2-\gamma\mu_\varphi)}{(1-\gamma\mu_\varphi)^2}, & \text{if } \gamma L_\varphi \in (0, \overline{\gamma L}_1\big(\frac{\mu_\varphi}{L_\varphi}\big)]; \\ \frac{\gamma(2-\gamma L_\varphi)}{(1-\gamma L_\varphi)^2}, & \text{if } \ \gamma L_\varphi \in (\overline{\gamma L}_1\big(\frac{\mu_\varphi}{L_\varphi}\big), 2). \end{cases}$$

Proposition 4.6.3 is derived by substituting the curvature expressions from the DCA splitting, as outlined in Proposition 4.6.1, into Theorem 4.2.1.

Corollary 4.6.1 (PGD sublinear rates for the residual gradient)**.** *Let $\varphi \in \mathcal{F}_{\mu_\varphi, L_\varphi}$ and $h \in \mathcal{F}_{0,\infty}$, with φ smooth (weakly) convex such that $\mu_\varphi \leq 0 < L_\varphi < \infty$. Consider N iterations of* (PGD) *with stepsize $\gamma L_\varphi \in \big(0, \overline{\gamma L}_1\big(\frac{\mu_\varphi}{L_\varphi}\big)\big]$, starting from x^0. Then:*

$$\tfrac{1}{2} \min_{0 \leq k \leq N}\{\|\nabla\varphi(x^k) + g_h^k\|^2\} \leq \frac{F(x^0) - F_{lo}}{\frac{\gamma(2-\gamma\mu_\varphi)}{(1-\gamma\mu_\varphi)^2} N},$$

where $g_h^k \in \partial h(x^k)$ for all $k = 1, \ldots, N$.

Corollary 4.6.1 follows immediately from Corollary 4.2.2 and corresponds to regime p_1, proved as being tight. The factor in the denominator is different that what we obtained for gradient descent for weakly convex functions. However, it is exactly the decrease after one iteration for (strongly) convex functions, as it can be equivalently written as $\frac{1}{\mu_\varphi}\big[-1 + (1-\gamma\mu_\varphi)^{-1}\big]$ when $\mu_\varphi \neq 0$.

Further on, we provide the performance bounds for PGD corresponding to the DCA results established for residual gradient and more than one iteration in Section 4.2.3. These bounds are obtained after performing the substitutions $\frac{L_2}{\mu_1} = 1 - \gamma\mu_\varphi$, $\frac{\mu_2}{\mu_1} = 1 - \gamma L_\varphi$ and $\mu_1 = \gamma^{-1}$. Conjecture 4.6.1 and Conjecture 4.6.2 are the translations in PGD terms of Conjecture 4.2.3 and Conjecture 4.2.4, respectively.

Conjecture 4.6.1 (Tight PGD rates for the residual gradient for F (weakly) convex)**.** *Let $\varphi \in \mathcal{F}_{\mu_\varphi, L_\varphi}$ and $h \in \mathcal{F}_{0,\infty}$, with φ smooth (weakly) convex such* *that* $\mu_\varphi \leq 0 < L_\varphi < \infty$. *Consider* N *iterations of* (PGD) *with stepsize* $\gamma L_\varphi \in \big(\overline{\gamma L}_1\big(\frac{\mu_\varphi}{L_\varphi}\big), 2\big)$, *starting from* x^0. *Then:*

$$\frac{1}{2}\min_{0\leq k\leq N}\{\|\nabla\varphi(x^k)+g_h^k\|^2\} \leq \frac{F(x^0)-F_{lo}}{\gamma\min\left\{P_N\left(1-\gamma\mu_\varphi, 1-\gamma L_\varphi\right)\,,\, E_N\left(1-\gamma L_\varphi\right)\right\}},$$

where $P_N(\eta,\rho) := \frac{(1+\eta)(1+\rho)}{\eta+\rho}\Big[N + \frac{(1-\eta)(1-\rho)}{\eta-\rho}\sum_{k=1}^{N}\left[E_N(\eta)-E_N(\rho)\right]_+\Big]$ *and* $g_h^k \in \partial h(x^k)$ *for all* $k = 1,\ldots,N$.

Conjecture 4.6.2 (Tight PGD rates for the residual gradient for F (strongly) convex)**.** *Let* $\varphi \in \mathcal{F}_{\mu_\varphi,L_\varphi}$ *and* $h \in \mathcal{F}_{0,\infty}$, *with* φ *smooth (strongly) convex such that* $0 \leq \mu_\varphi < L_\varphi < \infty$. *Consider* N *iterations of* (PGD) *with stepsize* $\gamma L_\varphi \in \big(0, 2\big)$, *starting from* x^0. *Then:*

$$\frac{1}{2}\min_{0\leq k\leq N}\{\|\nabla\varphi(x^k)+g_h^k\|^2\} \leq \frac{F(x^0)-F_{lo}}{\gamma\min\big\{E_N\left(1-\gamma\mu_\varphi\right)\,,\, E_N\left(1-\gamma L_\varphi\right)\big\}},$$

where $g_h^k \in \partial h(x^k)$ *for all* $k = 1,\ldots,N$.

Conjecture 4.6.2 can be proved in the cases with $\gamma L_\varphi \in \big(0, \overline{\gamma L}_1\big(\frac{\mu_\varphi}{L_\varphi}\big)\big]$ (using Theorem 4.2.2, corresponding to regime p_4) and $\gamma L_\varphi \in \big[\frac{2L_\varphi}{L_\varphi+\mu_\varphi}, 2\big)$ (using Theorem 4.2.3, corresponding to a subdomain of regime p_5). The rates are the same as in the unconstrained case and served as the initial base of formulating the conjectures.

Corollary 4.6.2 (Proved PGD rates for residual gradient for strongly convex F (p_4, subdomain of p_5))**.** *Let* $\varphi \in \mathcal{F}_{\mu_\varphi,L_\varphi}$ *and* $h \in \mathcal{F}_{0,\infty}$, *with* φ *smooth (strongly) convex, with* $\mu_\varphi \geq 0$. *Consider* N *iterations of* (PGD) *starting from* x^0 *with stepsize* $\gamma > 0$. *Then the following bounds hold:*

i) (Regime p_4*) If* $\gamma L_\varphi \in (0, \overline{\gamma L}_1\big(\frac{\mu_\varphi}{L_\varphi}\big)]$:

$$\tfrac{1}{2}\min_{0\leq k\leq N}\{\|\nabla\varphi(x^k)+g_h^k\|^2\} \leq \frac{F(x^0)-F_{lo}}{\gamma E_N(1-\gamma\mu_\varphi)}.$$

ii) (Regime p_5*) If* $\gamma L_\varphi \in \big[\frac{2L_\varphi}{L_\varphi+\mu_\varphi}, 2\big)$:

$$\tfrac{1}{2}\min_{0\leq k\leq N}\{\|\nabla\varphi(x^k)+g_h^k\|^2\} \leq \frac{F(x^0)-F_{lo}}{\gamma E_N(1-\gamma L_\varphi)}.$$

The unproved parts correspond to the challenging stepsize range $\gamma L_\varphi \in \big(\overline{\gamma L}_1\big(\frac{\mu_\varphi}{L_\varphi}\big), \overline{\gamma L}_N\big(\frac{\mu_\varphi}{L_\varphi}\big)\big)$, which, for gradient descent, requires distance-2 inequalities analysis. This final extension from GD to PGD is left as further work.

4.6.3 Convergence rates in gradient mapping

For this set of rates, we sometimes relax the convexity assumption on h, but preserving its smoothness, and only require $\mu_1 \geq 0$, hence $\gamma^{-1} + \mu_h \geq 0$, with strict inequality for stepsizes larger than $\frac{1}{L_\varphi}$.

Then regime r_1 corresponds to taking stepsizes $\gamma \leq \frac{1}{L_\varphi}$ for a splitting with φ and F nonconvex. Corollary 4.6.3 below translates the sublinear rates from Theorem 4.3.1 written for $\mu_2 \geq 0$, implying stepsizes $\gamma L_\varphi \in (0, 1]$.

Corollary 4.6.3 (PGD sublinear rates for the gradient mapping for weakly convex F ($\gamma L_\varphi \leq 1$, regime r_1))**.** *Let $\varphi \in \mathcal{F}_{\mu_\varphi, L_\varphi}$ and $h \in \mathcal{F}_{\mu_h, \infty}$, with φ smooth (weakly) convex such that $\mu_\varphi \leq 0 < L_\varphi < \infty$ and h nonsmooth, with $\mu_h \in (-L_\varphi, -\mu_\varphi)$. Consider N iterations of* (PGD) *with stepsize $\gamma L_\varphi \in (0, 1]$, starting from x^0. Then*

$$\tfrac{1}{2} \min_{0 \leq k \leq N} \{\|x^k - x^{k+1}\|^2\} \leq \frac{F(x^0) - F_{lo}}{\frac{2}{\gamma} + \mu_h - L_\varphi + \left[\frac{2}{\gamma} + \mu_h - L_\varphi + \frac{(L_\varphi + \mu_h)^2}{L_\varphi - \mu_\varphi}\right] N}.$$

The condition $\mu_h + L_\varphi \geq 0$ ensures in DCA the curvatures' relation $\mu_1 \geq \mu_2 \geq 0$ corresponding to regime r_1. This allows some degree of weak convexity even in the nonsmooth term. Corollary 4.6.3 has a direct correspondent in gradient descent for weakly convex functions with stepsizes lower than one, by taking $h = 0$ (hence $\mu_h = 0$) and using $x^k - x^{k+1} = \gamma \nabla \varphi_1(x^k)$. The shift in the denominator is usually obtained by taking $\gamma = \frac{1}{L_\varphi}$ (in DCA, this means $\mu_2 = 0$), as we had in the case of gradient descent.

Within the regimes r_3 and r_4, the stepsize γ lies in the interval $[\frac{1}{L_\varphi}, \bar{\gamma}^1(L_\varphi, \mu_\varphi, \mu_h)]$. Regime r_3 holds when φ and F are nonconvex. When $\mu_h = 0$, then the threshold expression reduces to $\overline{\gamma L}_1(\kappa)$ defined for gradient descent, with $\kappa = \frac{\mu_\varphi}{L_\varphi}$.

Regime r_4 holds when F is strongly convex and $\mu_\varphi + \mu_h \geq 0$; in particular, this includes the case of φ being strongly convex. For stepsizes $\gamma \in (\bar{\gamma}^1(L_\varphi, \mu_\varphi, \mu_h), \frac{2}{L_\varphi + \mu_\varphi})$, regime r_5 is active if F is nonconcave.

Conjecture 4.6.3 is very close in spirit to Conjectures 4.6.1 and 4.6.2, which are stated for the residual subgradient. The main difference is that those conjectures explicitly impose the standard assumption $\mu_h = 0$ and $L_h = \infty$ (i.e., h is convex), whereas here we do not. This statement translates Conjecture 4.3.1 from the DCA setting to the PGD setting.

Conjecture 4.6.3 (Tight PGD rates for the gradient mapping for general F)**.** *Let $\varphi \in \mathcal{F}_{\mu_\varphi, L_\varphi}$ and $h \in \mathcal{F}_{\mu_h, L_h}$. Consider $N + 1 \geq 2$ iterations of* (PGD) *with stepsize $\gamma \in \left(0, \min\left\{\frac{2}{L_\varphi - \mu_h}, \frac{1}{[\mu_h]_-}\right\}\right)$, starting from x^0. The following hold:*

i) (F nonconvex-nonconcave) If $\mu_\varphi + \mu_h \le 0$*,* $L_\varphi + L_h > 0$ *and* $\gamma > \frac{1}{L_\varphi}$*, then:*

$$\frac{\gamma^{-1} + \mu_h}{2} \min_{0 \le k \le N} \left\{\|x^k - x^{k+1}\|^2\right\} \le$$

$$\frac{F(x^0) - F(x^{N+1})}{(\frac{2}{\gamma} + \mu_h - L_\varphi) + \min\left\{P_N\left(1 - \gamma\mu_\varphi, 1 - \gamma L_\varphi\right),\, E_N\left(1 - \gamma L_\varphi\right)\right\}},$$

where $P_N(\eta, \rho) := \frac{(1+\eta)(1+\rho)}{\eta+\rho}\left(N + \frac{(1-\eta)(1-\rho)}{\eta-\rho}\sum_{k=0}^N [E_k(\eta) - E_k(\rho)]_+\right)$.

ii) (F (strongly) convex) If $\mu_\varphi + \mu_h \ge 0$*, then:*

$$\frac{\gamma^{-1} + \mu_h}{2}\|x^k - x^{k+1}\|^2 \le$$

$$\frac{F(x^0) - F(x^{N+1})}{(\frac{2}{\gamma} + \mu_h - L_\varphi) + \min\left\{E_N\left(1 - \gamma\mu_\varphi\right),\, E_N\left(1 - \gamma L_\varphi\right)\right\}}.$$

iii) (F (strongly) concave) If $L_\varphi + L_h \le 0$ *and* $\gamma \in \left(0,\, \min\left\{\frac{1}{[\mu_h]_-}, \frac{1}{L_\varphi}\right\}\right)$*, then:*

$$\frac{\gamma^{-1} - L_\varphi}{2}\|x^k - x^{k+1}\|^2 \le$$

$$\frac{F(x^0) - F(x^{N+1})}{(\frac{2}{\gamma} + \mu_h - L_\varphi) + \min\left\{E_N\left(\frac{1+\gamma L_h}{1-\gamma L_\varphi}\right),\, E_N\left(\frac{1+\gamma\mu_h}{1-\gamma L_\varphi}\right).\right\}}.$$

Further on, we state corollaries obtained by translating the proved regimes for the gradient mapping analysis of DCA from Section 4.3.

Corollary 4.6.4 (PGD rates for gradient mapping (regimes r_3, r_4))**.** *Let* $\varphi \in \mathcal{F}_{\mu_\varphi, L_\varphi}$ *and* $h \in \mathcal{F}_{\mu_h, \infty}$*, with* φ *smooth such that* $\mu_\varphi < L_\varphi < \infty$ *and* h *nonsmooth, with* $|\mu_h| < L_\varphi$*. Consider* N *iterations of* (PGD) *starting from* x^0 *with stepsize* $\gamma \in (0, \overline{\gamma}^1(L_\varphi, \mu_\varphi, \mu_h)]$*, where the threshold is defined in Definition 4.6.1. Then the following bounds hold:*

i) (F (weakly) convex, regime r_3*) If* $\mu_h + \mu_\varphi \le 0$ *and* $\gamma \in [\frac{1}{L_\varphi}, \overline{\gamma}^1(L_\varphi, \mu_\varphi, \mu_h)]$*:*

$$\tfrac{1}{2}\min_{0 \le k \le N}\{\|x^k - x^{k+1}\|^2\} \le \frac{F(x^0) - F_{lo}}{\frac{2}{\gamma} + \mu_h - L_\varphi + \left[\frac{(\mu_h - L_\varphi + \frac{2}{\gamma})(\mu_h - \mu_\varphi + \frac{2}{\gamma})}{L_\varphi + \mu_\varphi - \frac{2}{\gamma}}\right]N}.$$

ii) (F (strongly) convex, regime r_4) If $\mu_h + \mu_\varphi \geq 0$ and $\gamma \in (0, \bar{\gamma}^1(L_\varphi, \mu_\varphi, \mu_h)]$:

$$\frac{1}{2} \min_{0 \leq k \leq N} \{\|x^k - x^{k+1}\|^2\} \leq \frac{F(x^0) - F_{lo}}{\frac{2}{\gamma} + \mu_h - L_\varphi + \frac{-1 + \left(\frac{1-\gamma\mu_\varphi}{1+\gamma\mu_h}\right)^{-2N}}{1 - \left(\frac{1-\gamma\mu_\varphi}{1+\gamma\mu_h}\right)}}.$$

Regimes r_2 and r_6, respectively, correspond to using stepsizes $\gamma \leq \frac{1}{L_\varphi}$ for φ concave and F nonconcave or concave, respectively. We give in Corollary 4.6.5 the rate for the F nonconcave, implied by h nonsmooth ($L_h = \infty$).

Corollary 4.6.5 (PGD sublinear rate for gradient mapping for nonconcave F ($\gamma L_\varphi \leq 1$, regime r_2))**.** *Let $\varphi \in \mathcal{F}_{\mu_\varphi, L_\varphi}$ and $h \in \mathcal{F}_{\mu_h, \infty}$, with φ smooth (weakly) convex such that $\mu_\varphi \leq 0 < L_\varphi < \infty$ and h nonsmooth, with $\mu_h \leq -L_\varphi$. Consider N iterations of* (PGD) *with stepsize $\gamma L_\varphi \in (0, 1]$, starting from x^0. Then the following bound holds*

$$\frac{1}{2} \min_{0 \leq k \leq N} \{\|x^k - x^{k+1}\|^2\} \leq \frac{F(x^0) - F_{lo}}{(\frac{2}{\gamma} + \mu_h - L_\varphi)(N+1)}.$$

Lemma 4.6.1 provides an example of descent lemma translation between PGD and DCA, corresponding to regime r_3 from Lemma 4.5.6, assuming the standard setup of PGD with $\mu_h = 0$ and $L_h = \infty$, and stepsizes $\gamma L_\varphi \in [1, 2)$. It leads to proving the first case of Corollary 4.6.4.

Lemma 4.6.1 (PGD one-step descent for stepsizes $\gamma L_\varphi \in [1, 2)$ (regime r_3))**.** *Let $\varphi \in \mathcal{F}_{\mu_\varphi, L_\varphi}$, $h \in \mathcal{F}_{0,\infty}$ and $\gamma \in [\frac{1}{L_\varphi}, \frac{2}{L_\varphi})$. Then after two iterations of PGD with stepsize γ connecting x^k, x^{k+1} and x^{k+2} it holds*

$$\begin{aligned} F(x^k) - F(x^{k+1}) \geq & \frac{[(2 - L_\varphi)(2 - \mu_\varphi) - 1]}{\gamma(2 - L_\varphi - \mu_\varphi)} \frac{1}{2}\|x^k - x^{k+1}\|^2 + \\ & \frac{1}{\gamma(2 - L_\varphi - \mu_\varphi)} \frac{1}{2}\|x^{k+1} - x^{k+2}\|^2. \end{aligned} \tag{4.28}$$

Lemma 4.6.1 is the extension to PGD of the similar result obtained for GD in Lemma 3.5.2, after replacing the GD iteration $x^{k+1} = x^k - \gamma \nabla \varphi(x^k)$. Therein, it is involved in deriving a stepsize schedule gradually increasing from $\bar{\gamma}^1(L_\varphi, \mu_\varphi, 0)$ towards $\frac{2}{L_\varphi + \mu_\varphi}$ (if φ is (strongly) convex) or up to certain stepsize limit in the nonconvex case; in both cases, this procedure gives a better worst-case guarantee than the best constant stepsize (see details in Section 3.3.2). Based on Lemma 4.6.1, the same schedule gives exactly the same performance of gradient mapping for PGD. In Section 4.7.2, we show how this schedule can be translated to DCA.

Remark 4.6.2 (Recovering GD results from gradient mapping)**.** *Recall that the gradient mapping is defined by*

$$\gamma^{-1}(x^k - x^{k+1}) = \nabla\varphi(x^k) + g_h^{k+1}, \quad g_h^{k+1} \in \partial h(x^{k+1}),$$

*and note that GD is recovered by taking $h = 0$. Hence, starting from DCA/PGD lemmas stated in terms of iterate progress (or the gradient mapping) and substituting these differences through $\nabla\varphi(x^k)$, one would recover **all** results obtained for GD. However, in the DCA/PGD setting we only obtain descent results based on distance-1 inequalities, and not the more intricate regimes that rely on distance-2 inequalities. Ideally, from a didactic perspective, the DCA lemmas would provide a unifying viewpoint.*

4.6.4 Convergence rates for proximal point method

We conclude this section by providing the performance bounds for the proximal point method (PPA) using variable stepsizes. Recall that PPA belongs to the PGD setting with $\varphi = 0$ and $h \in \mathcal{F}_{0,\infty}$. The PPA iteration with stepsize γ_k reads

$$x^{k+1} = \text{prox}_{\gamma_k h}(x^k). \tag{4.29}$$

Proposition 4.6.4 (Sublinear rates for PPA in residual gradient and gradient mapping)**.** *Let $h \in \mathcal{F}_{0,\infty}$ and consider N iterations of (4.29) starting from x^0, using stepsizes $\{\gamma_k\}$, with $k = 0, \ldots, N-1$. Then the following bounds hold:*

$$\textit{residual gradient:} \quad \min_{0\le k\le N}\{\|g^k\|^2\} \le \frac{h(x^0) - h^*}{\sum_{k=0}^{N-1}\gamma_k},$$

where $g^k \in \partial h(x^k)$, and

$$\textit{gradient mapping:} \quad \min_{0\le k\le N}\{\|x^k - x^{k+1}\|^2\} \le \frac{h(x^0) - h^*}{2\gamma_{N-1}^{-1} + \sum_{k=0}^{N-2}\gamma_k^{-1}}.$$

Proof. From the PGD-DCA equivalence, in DCA terms we get $\mu_1 = \mu_2 = L_2 = (\gamma_k)^{-1}$ and $L_1 = \infty$ at each iteration k. This setting corresponds to regime p_1 for both metrics (residual gradient and gradient mapping) and to descent inequalities

$$h(x^k) - h(x^{k+1}) \ge \gamma_k \|g_h^k\|^2$$

and

$$h(x^k) - h(x^{k+1}) \ge \frac{1}{\gamma_k}\|x^k - x^{k+1}\|^2,$$

respectively. For residual gradient, the conclusion follows by telescoping and taking the minimum norm. In the second case, one can check that this inequality is exactly the sufficient decrease inequality from Proposition 4.1.2, with $\mu_1 = \mu_2 = \gamma_k^{-1}$, so we additionally have that $h(x^k) - h^* \geq h(x^k) - h(x^{k+1}) \geq \gamma_{N-1}^{-1}\|x^k - x^{k+1}\|^2$. Adding this inequality to the telescoped summation and taking the minimum norm gives the conclusion. □

The result in Proposition 4.6.4 is complementary to the tight rates from [113, Theorem 4.1] and [113, Conjecture 4.2], on the measures $\frac{h(x^N)-h^*}{\|x^0-x^*\|^2}$ and $\frac{g^N}{\|x^0-x^*\|}$, where $g^N \in \partial h(x^N)$, respectively.

4.7 Curvature shifting technique

We assume that the convex conjugate of f_1 and of any curvature adjustment $f_1 - \lambda\frac{\|\cdot\|^2}{2}$, where $\lambda \leq \mu_1$, can be computed efficiently. This leads to the natural question: given a splitting $F = f_1 - f_2$, what is the optimal curvature shift λ in the decomposition $F = (f_1 - \lambda\frac{\|\cdot\|^2}{2}) - (f_2 - \lambda\frac{\|\cdot\|^2}{2})$? In fact, this is a standard approach for addressing weak convexity in the function f_2, by lifting it to a convex function, with some $\lambda \leq \mu_2 < 0$, and then applying the DCA iteratively.

We show that this approach is suboptimal: the sublinear rates constants, measured via the largest denominator p_i, can be improved. Specifically, the constant in sublinear rates for a nonconvex-nonconcave objective F (where $\max\{\mu_1, \mu_2\} < \min\{L_1, L_2\}$) is optimized with respect to the curvature shift λ by maximizing the denominator. In some cases, as indicated in the *Ratio* column from Table 4.7.1, the improvement is significant.

Let $\tilde{f}_i^\lambda := f_i - \lambda\frac{\|\cdot\|^2}{2}$, with $i = \{1, 2\}$, be the curvature adjusted functions, and $P_\lambda := P(L_1, \mu_1, L_2, \mu_2, \lambda)$ be the denominator corresponding to one of the six possible regimes p_i, determined by the initial splitting curvatures and the parameter λ. Initially, $\lambda = 0$ and the denominator is P_0. In Table 4.7.1 we use p_i, along with its value, to represent the regime before and after the curvature adjustment. Given the analytical expressions for all six regimes, we can easily numerically compute $\lambda^* = \arg\max_\lambda P_\lambda$.

The examples in Table 4.7.1 suggest the following observations. First, consider both functions to be (strongly) convex. If $\mu_1 \leq \mu_2$, then the best splitting is achieved by making $\tilde{f}_1$ convex, hence $\lambda = \mu_1$. Notably, when $\mu_1 = \mu_2$, both functions become convex, which surprisingly implies that the initial strong convexity may actually slow down the algorithm. If $\mu_1 > \mu_2$, the optimal splitting occurs when f_2 is shifted to a weakly convex function. Furthermore,

Table 4.7.1: Improvement of DCA convergence rates for nonconvex-nonconcave objectives F ($\max\{\mu_1, \mu_2\} < \min\{L_1, L_2\}$) using curvature shifting with optimal λ^*. Relative improvement (*Ratio*) is defined as $\frac{P(\lambda^*)-P(0)}{P(0)}$, where $P(0)$ is the initial denominator before splitting and $P(\lambda^*)$ is the optimal (largest) one.

Setup	μ_1	L_1	μ_2	L_2	$P(0)$	λ^*	$P(\lambda^*)$	**Ratio**
$\mu_1 > \mu_2$	0.2	3	0.1	4	$p_2 = 0.4535$	0.1221	$p_3 = 0.6031$	33%
	0.2	1000	0.1	3	$p_1 = 0.3255$	0.1494	$p_1 = p_3 = 0.358$	10%
	1	2	0.5	1.5	$p_1 = 0.7600$	0.6733	$p_3 = 2.0321$	167%
$\mu_1 = \mu_2$	1	4	1	3	$p_1 = 0.4583$	μ_1	$p_1 = p_2 = 0.8333$	81%
$\mu_1 < \mu_2$	0.1	3	0.2	4	$p_2 = 0.4539$	μ_1	$p_2 = 0.602$	32%
	0.001	4	0.002	3	$p_1 = 0.4531$	μ_1	$p_1 = 0.5835$	28%
$\mu_1 > 0$	2	4	-1.75	3	$p_5 = 0.4688$	-0.4855	$p_3 = 0.52$	11%
$\mu_1 + \mu_2 > 0$	2.99	4	-2.9	3	$p_5 = 0.4994$	-0.935	$p_5 = 0.5076$	1.6%
$\mu_1 > 0$ $\mu_1 + \mu_2 < 0$	1	2	-1.5	1.5	–	-0.6526	$p_3 = 0.8516$	–

even when starting with f_2 as weakly convex, the best curvature maintains this weak convexity. Additionally, in case of a *bad* decomposition where $\mu_1 + \mu_2 < 0$ (providing no convergence guarantee for DCA iterations), an appropriate λ can ensure convergence. In fact, the optimal denominator is reached for some $\mu_2 < 0$.

This collection of seemingly surprising results involving weak convexity of f_2 is explained by the equivalence with the proximal gradient descent (PGD) (see Section 4.6) and its faster convergence when using a stepsize larger than the inverse Lipschitz constant.

Benefits on Smooth Functions. The objective F is smooth when both f_1 and f_2 are smooth, raising the question of why to not use gradient descent (GD) directly. We compare the rates of DCA, having the iteration $x^+ = \nabla f_1^*(\nabla f_2(x))$, to GD, whose iteration with stepsize $\gamma \in (0, \frac{2}{L_F})$ reads $x^+ = x - \gamma \nabla F(x)$. We use the criteria $\frac{1}{2}\min\{\|\nabla F(x)\|^2, \|\nabla F(x^+)\|^2\} \leq \frac{F(x)-F(x^+)}{p}$, where p denotes the worst-case denominator for DCA or GD. In Example 4.7.1 we show that, for smooth nonconvex functions, the optimal rate for DCA applied can surpass the optimal rate of GD.

Example 4.7.1 (Benefits of curvature shifting for smooth-nonconvex F)**.** *Let* $L_F = -\mu_F = 1$, *and the decomposition* $F = f_1 - f_2$, *with parameters* $\mu_1 = 1.5$, $L_1 = 2$, $\mu_2 = 1$ *and* $L_2 = 2.5$. *This corresponds to regime* p_2, *with* $p_2 = 0.9167$. *The optimal stepsize for GD,* $\gamma^* = \frac{2}{\sqrt{3}}$ *(Abbaszadehpeivasti et al. [2, Theorem 3]), yields the denominator* $p_{GD} = 1.5396$, *which is* 67% *larger than for the initial DCA splitting. However, the best DCA splitting achieved by subtracting* *curvature* $\lambda^* = 1.0091$*, leading to* $F = \tilde{f}_1^\lambda - \tilde{f}_2^\lambda$ *with* $\tilde{f}_1^\lambda \in \mathcal{F}_{0.4009,0.9009}$ *and* $\tilde{f}_2^\lambda \in \mathcal{F}_{-0.0991,1.4009}$*, corresponds to regime* p_3*, with* $p_{DCA^*} = 1.724$*, which is* 12% *larger than* p_{GD}*. Thus, with appropriately chosen curvatures in DCA, we can improve convergence rates even in the smooth case.*

A different perspective is to assume as given the curvature bounds of the nonconvex-nonconcave objective function F and map the optimized procedure depending on two of the four curvatures. Specifically, in Figure 4.7.1 we provide an example with $\mu_F = -0.5$ and $L_F = 1.5$, and illustrate all possible regimes after one iteration. Note that $\mu_F = \mu_1 - L_2$ and $L_F = L_1 - \mu_2$; further on, we examine the regimes based on the ranges of L_2 and μ_2. The condition $\mu_1 \geq 0$ implies that $L_2 = \mu_1 - \mu_F \geq -\mu_F$, while the condition $\mu_1 + \mu_2 > 0$ that $\mu_2 > -\mu_1 = L_2 + \mu_F$. The contour lines represent values of denominators p_i, with $i = \{1, 2, 3, 5\}$. The *red* points mark the initial curvature values of L_2 and μ_2, whereas the *green* dots indicate the points with the largest possible p_i obtained through the optimal choice of λ. Since these shifts are linear in λ, the dashed lines connecting the dots have a slope of one.

For example, in the case where $L_2 = 1$ and $\mu_2 = 0.75$, we have $\mu_1 < \mu_2$ and the best splitting is obtained by shifting to the lowest possible value of L_2, corresponding to $\mu_1 = 0$. In all other examples, $\mu_1 > \mu_2$ and the optimal splittings are found within regime p_3, where $\mu_2 < 0$.

4.7.1 Numerical experiments: Sparse PCA

We provide numerical evidence supporting the efficiency of curvature shifting, aligning with the optimal splittings proposed in Section 4.7, by exploring the sparse principal component analysis (SPCA) problem. The focus is on the (best) residual gradient metric. Following Journée et al. [63] and Themelis et al. [119], we include elastic-net regularization and solve

$$\underset{x \in \bar{B}(0,1)}{\text{minimize}}\, F(x) := \kappa\|x\|_1 + \eta\frac{\|x\|^2}{2} - \frac{1}{2}x^\top \Sigma x,$$

where $\bar{B}(0,1) := \{x \mid \|x\|_2 \leq 1\}$ is the closed Euclidean unit ball, $\Sigma = A^\top A$ the sample covariance matrix, and κ, η are the l_1-(sparsity inducing) and l_2-regularization parameters, respectively. We denote

$$f_1(x) := \kappa\|x\|_1 + \eta\frac{\|x\|^2}{2} + \delta_{\bar{B}(0,1)}(x),$$

$$f_2(x) := \frac{1}{2}x^\top \Sigma x,$$

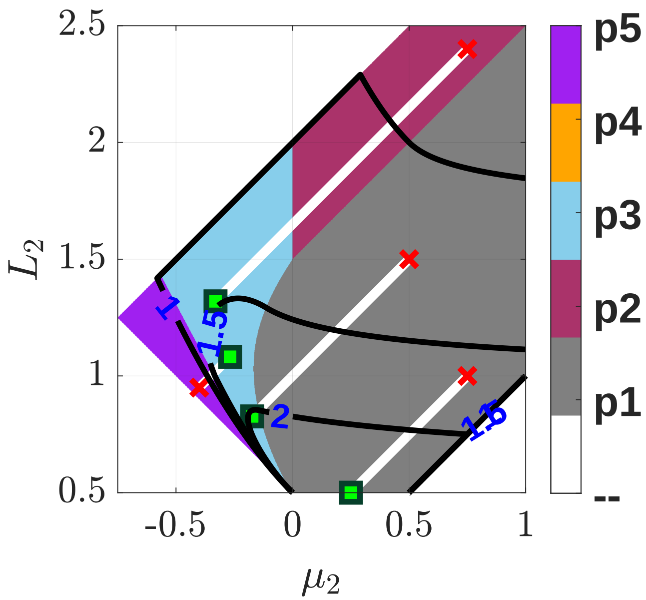

Figure 4.7.1: All regimes for a fixed objective function F with $\mu_F = -0.5$ and $L_F = 1.5$, along with several mappings of the optimal splittings, shown as transitions from *red* to *green* dots along dashed lines with a slope of 1.

with $\delta_C(x)$ as the indicator function of a set C. Note that $\mu_1 = \eta$, $L_1 = \infty$, $\mu_2 = \min\{\Lambda(\Sigma)\}$ and $L_2 = \max\{\Lambda(\Sigma)\}$, where $\Lambda(\Sigma)$ denotes the eigenvalues of Σ. Further, we consider the splittings $\tilde{f}_i^{\lambda}(x) \coloneqq f_i(x) - \lambda\frac{\|x\|^2}{2}$, with $i = \{1, 2\}$, parametrized by λ (curvature shift); note that $F(x)$ remains unchanged. The subdifferential of the convex conjugate of $\tilde{f}_1^{\lambda}$ can be computed in closed form for any $\lambda \leq \eta$ (see Section 4.B.1). We generated a sparse random matrix $A \in \mathbb{R}^{20n\times n}$ $(n = 200)$, with 10% density, and normalized $\Sigma = AA^\top$ by its maximum eigenvalue, obtaining $f_2 \in \mathcal{F}_{\mu_2=0.3882, L_2=1}$. Parameters κ and η were selected to ensure a sparse but non-trivial solution, and our focus is to demonstrate the impact of curvature shifting on convergence.

We use $M = 1000$ random initial points within the unit ball, and only keep the $\bar{M}$ runs that converge to the same solution (non-trivial and with desired sparsity). We compare convergence rates for various splittings, namely $\{0, \pm\lambda^*, \pm 0.5\lambda^*, \lambda_{\max}\}$, where λ^* is the optimal curvature shift defined in Section 4.7 and $\lambda_{\max} \coloneqq \frac{\mu_1+\min\{\mu_1,\mu_2\}}{2}$ is the maximum shift guaranteeing convergence. We report in Table 4.7.2 the average number of iterations N_ε required to reach a specific accuracy level $\varepsilon \in \{10^{-2}, 10^{-4}, \ldots, 10^{-10}\}$ of the squared (sub)gradient norm of $F(x^k)$ (which is equal to $\arg\min_{0\le k\le N} \|\tilde{g}_1^k - \nabla \tilde{f}_2^\lambda(x^k)\|^2$), for $\varepsilon = \{10^{-1}, 10^{-2}, \dots, 10^{-12}\}$ and $\tilde{g}_1^k \in \partial \tilde{f}_1^\lambda(x^k)$.

We consider two cases: (i) $\eta = \mu_1 > \mu_2$ and (ii) $\eta = \mu_1 < \mu_2$. Complete tables with the experimental results are provided in Section 4.B.

Case 1: $\eta > \mu_2$. Parameters: $\eta = \mu_1 = 0.5$, $\kappa = 0.02$ (400 runs kept out of 1000). Here, $\lambda^* = 0.4413$ and the maximum *theoretical* shift is $\lambda_{\max} = \frac{\mu_1+\mu_2}{2} = 0.4441$. With $\lambda > \mu_2 = 0.3882$, $\tilde{f}_2^\lambda$ becomes weakly convex. The results are reported in Table 4.B.1 (in the appendix) and Figure 4.7.2. Using λ^* achieves at least a twofold acceleration compared to $\lambda = 0$. Higher λ values, increasing the nonconvexity of $\tilde{f}_2^\lambda$, further improve the convergence. Conversely, adding curvature to both functions ($\lambda < 0$) slows the convergence. The improved convergence rates observed when $\tilde{f}_2^\lambda$ becomes weakly convex further highlight the significance of studying DCA for weakly convex functions.

Case 2: $\eta < \mu_2$. Parameters: $\eta = 0.2$, $\kappa = 0.02$ (703 runs kept out of 1000). The results are reported in Table 4.B.1 (in the appendix) and Figure 4.7.3. We get $\lambda^* = \mu_1 = 0.2$. Using λ^* to make $\tilde{f}_1^\lambda$ convex improves the convergence speed by approximately 20% compared to the initial splitting. Generally, decreasing the curvature of f_2 improves the convergence.

In conclusion, λ^* from worst-case analysis is also effective for practical curvature shifting, despite better-than-predicted performance. Reducing convexity in both functions improves the rates, similar to larger stepsizes enhancing PGD convergence.

4.7.2 Curvature shifting based schedules for DCA

Based on a single-iteration analysis, in the previous section we observed better performance when optimizing the progress after one iteration. A natural improvement suggestion of the procedure would be to employ the conjectured bounds for any number of iterations and compute the best curvature shift numerically.

In the setting with $\mu_1 > 0$, $L_1 = \infty$, $L_2 \in (0, \infty)$ and $\mu_2 < \mu_1$, there exists additional flexibility in choosing the shift. In this case, we can exploit the connection with PGD when the initial curvature choice of the problem allows it. More precisely, (optimized) stepsize schedules $\{\gamma_k(L_\varphi, \mu_\varphi)\}$ for PGD can be translated to DCA.

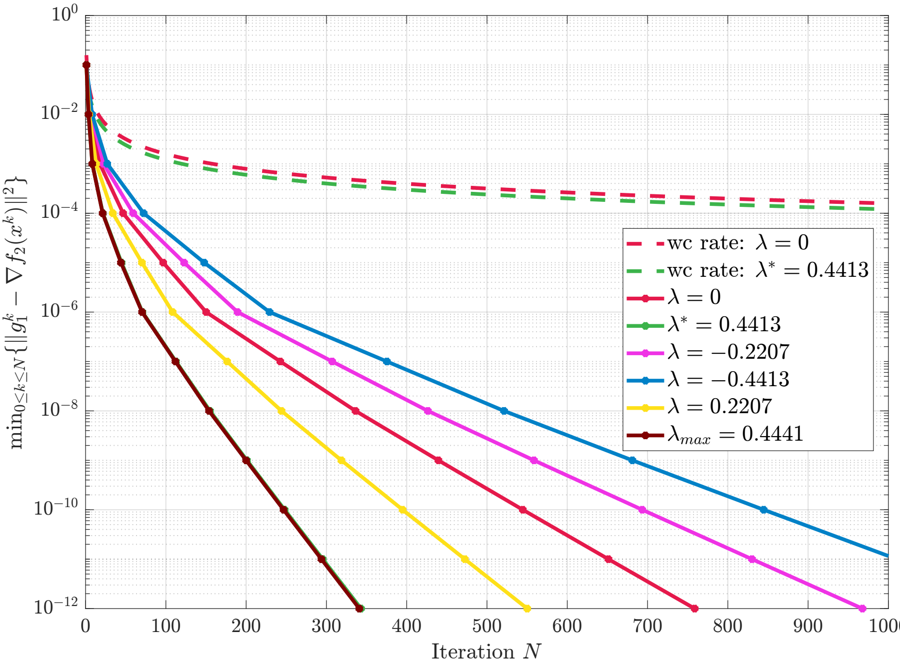

Figure 4.7.2: Curvature shifting experiment for $\mu_1 < \mu_2$ (**case 1**). Setting: $\eta = 0.5 > \mu_2$, $\kappa = 0.02$, $\lambda = \{0, \pm\lambda^*, \pm 0.5\lambda^*, \lambda_{\max}\}$. Lowest values are obtained for $\lambda_{\max} = 0.4441$. Average number of iterations N_ε to reach ε accuracy of the best subgradient norm $\|g_1^k - \nabla f_2(x^k)\|$, where $g_1^k \in \partial f_1(x^k)$. Although the actual rate exceeds theoretical guarantees (*wc rate*), there is correlation between optimizing the worst-case for λ and better performance in the experiment.

The shifted curvatures at iteration k are $(\mu_1^\lambda)_k = \mu_1 - \lambda_k$, $(L_1^\lambda)_k = \infty$ and $(\mu_2^\lambda)_k = \mu_2 - \lambda_k$, $(L_2^\lambda)_k = L_2 - \lambda_k$. The corresponding PGD setting, when taking the stepsize $\gamma = \frac{1}{(\mu_1^\lambda)_k}$, is characterized by the curvatures $\mu_h = 0$, $L_h = \infty$, $\mu_\varphi = \gamma^{-1} - L_2$ and $L_\varphi = \gamma^{-1} - L_2$. This setting corresponds to smooth minimization over convex set. Then we get

$$\{\lambda_k\} = \mu_1 - \{[\gamma_k(\mu_1 - L_2, \mu_1 - \mu_2)]^{-1}\} \tag{4.30}$$

as the optimized splitting schedule for DCA, adapting the curvature bounds of f_1^λ and f_2^λ at each iteration.

One such optimized schedule is the one provided in Section 3.3.2 for gradient descent, which we showed that also holds for PGD (see discussion after Lemma 4.6.1).

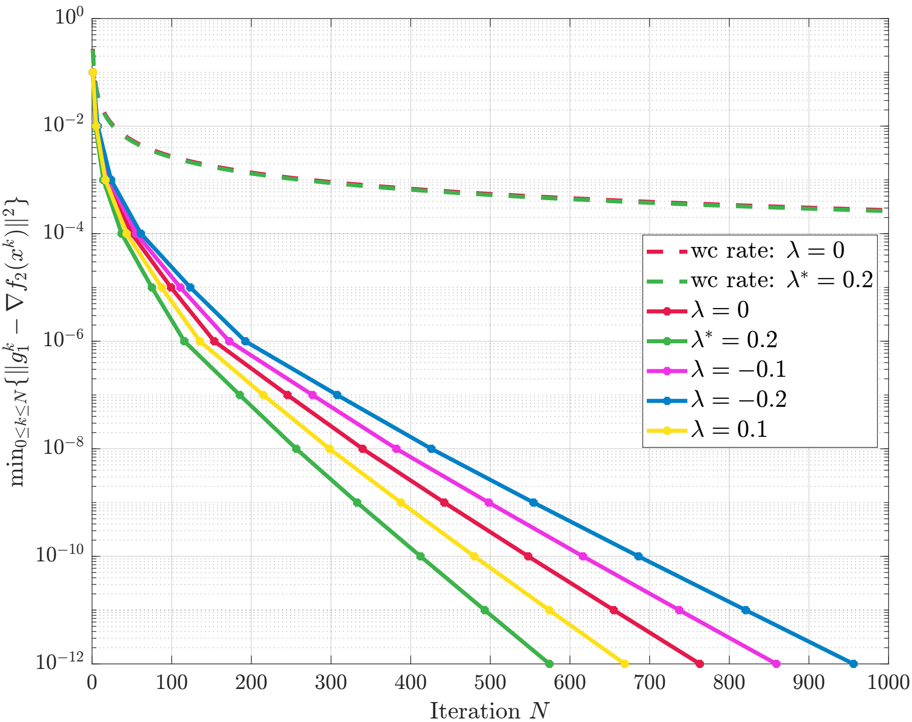

Figure 4.7.3: Curvature shifting experiment for $\mu_1 < \mu_2$ (**case 2**). Setting: $\eta = 0.2 = \mu_1 < \mu_2$, $\kappa = 0.02$, $\lambda = \{0, \pm\lambda^*, \pm 0.5\lambda^*\}$. Average number of iterations N_ε to reach ε accuracy of the best subgradient norm $\|g_1^k - \nabla f_2(x^k)\|$, where $g_1^k \in \partial f_1(x^k)$. Lowest values are obtained for $\lambda^* = \lambda_{\max} = \mu_1 = 0.2$. There is correlation between optimizing the worst-case for λ and better performance in the experiment.

Assessing the practical performance of such schedules for DCA is left as future work.

4.8 Conclusion

We thoroughly examined a single (DCA) iteration applied to the DC framework extended to accommodate one weakly convex function. We characterized six distinct regimes for the objective decrease based either on subgradient differences or iterate progress, and conjecture, based on numerical observations, that certain regimes remain tight across multiple iterations. Sublinear convergence rate

Table 4.7.2: Average number of iterations N_ε for $\kappa = 0.02$ and (Case 1) $\eta = 0.5 = \mu_1 > \mu_2$; (Case 2) $\eta = 0.2 = \mu_1 < \mu_2$, for $\{0, \pm\lambda^*, \pm 0.5\lambda^*, \lambda_{\max}\}$.

Case	λ \ ε	10^{-2}	10^{-4}	10^{-6}	10^{-8}	10^{-10}
1	0	5.74	46.70	150.17	336.27	544.63
	0.4413	3.65	21.49	70.75	154.93	247.93
	-0.2207	6.83	59.32	189.31	426.39	693.40
	-0.4413	8.05	72.30	229.13	521.24	844.69
	0.2207	4.54	33.91	108.43	244.05	394.92
	0.4441	3.63	21.23	70.38	153.78	246.61
2	0	5.88	49.85	153.46	339.67	547.61
	0.2	4.32	37.10	115.90	256.33	412.39
	-0.1	6.48	55.52	172.24	382.11	616.29
	-0.2	7.06	61.44	192.46	426.00	686.04
	0.1	5.26	43.47	135.39	298.14	479.93

results follow as a corollary of this analysis. For the non-tight regimes after one iteration, we are able to conjecture their exact behaviour, drawing inspiration from the rates for gradient descent developed in Chapter 3.

Although the rates for regimes p_1 and p_2 were stated in [3] without explicit descent lemmas, such lemmas could likely be extracted with a further refinement of the arguments therein. In contrast, for the gradient mapping criterion, even though convergence rates for regimes r_1 and r_2 were also obtained in [3] by applying Toland duality [122] to the residual gradient rates, the corresponding descent lemmas do not follow from the same approach. We therefore view these descent lemmas as a novel contribution to the convergence analysis of (standard) DCA.

The (DC) structure of the objective facilitates curvature shifting by adding/-subtracting $\frac{\lambda}{2}\|x\|^2$ to both terms. However, applying the (DCA) iteration on a split with one weakly convex function can yield better rates than using a modified split that achieves convexity for both functions, as shown by several examples and confirmed by numerical experiments.

We highlight the strong link between PGD and DCA, leveraging their iteration equivalence to translate DCA rates to PGD setups. This connection suggests the potential for an optimized splitting of the objective function used in DCA, possibly varying at each iteration, inspired by stepsize schedules used for (proximal) gradient descent. We leave the practical testing of those as future work.

Appendix

4.A Tightness in residual gradient

For each of the six regimes denoted p_i, with $i = 1, \dots, 6$, we provide numerical examples of the corresponding worst-case values, defined as $\mathrm{wc}(N) = \frac{1}{2}\min_{0\leq k\leq N}\{\|g_1^k - g_2^k\|^2\}$, where $g_1^k \in \partial f_1(x^k)$, $g_2^k \in \partial f_2(x^k)$ and the initial condition is fixed to $F(x^0) - F(x^N) \leq 1$. See Figure 4.A.1 for regime p_1 when f_2 is convex or weakly convex, Figure 4.A.2 for regimes p_2 and p_3, Figure 4.A.3 for regime p_5 with $L_2 > 0$ or $L_2 < 0$ and Figure 4.A.4 for regimes p_4 and p_6. The rates from Corollary 4.2.1 are predicted to be sublinear, with $\mathrm{wc}(N) = {}^1\!/_{p_i N}$. This holds for regimes p_1, p_2 and p_3, as stated in Proposition 4.2.1 and Conjecture 4.2.1. For the domains of regimes p_4, p_5 and p_6, however, the tight rates are obtained by a different analysis and are not sublinear for any N, as seen in Section 4.2.3.

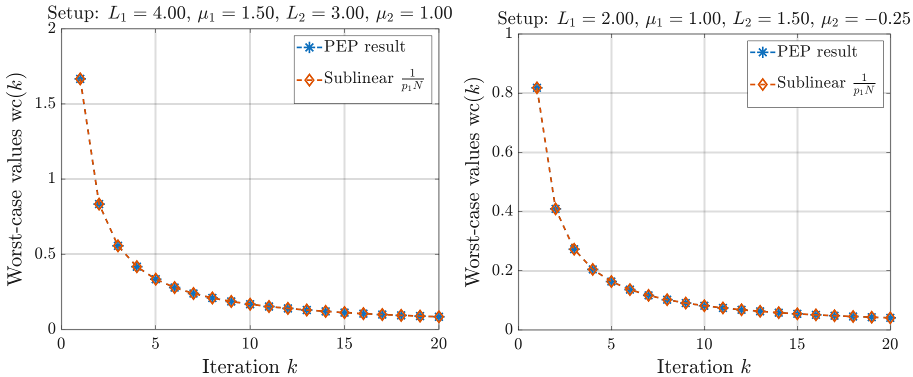

Figure 4.A.1: Examples for regime p_1 with f_2 convex (*left*) and f_2 weakly convex (*right*), showing exactness of our expressions.

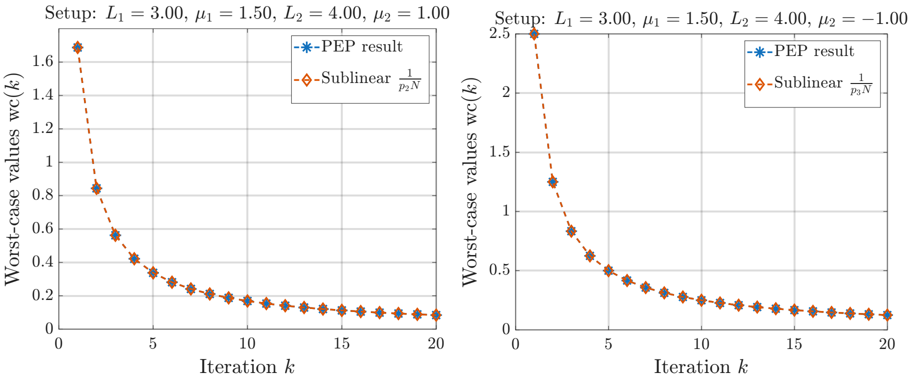

Figure 4.A.2: Examples for regimes p_2 (*left*) and p_3 (*right*), showing exactness of our expressions.

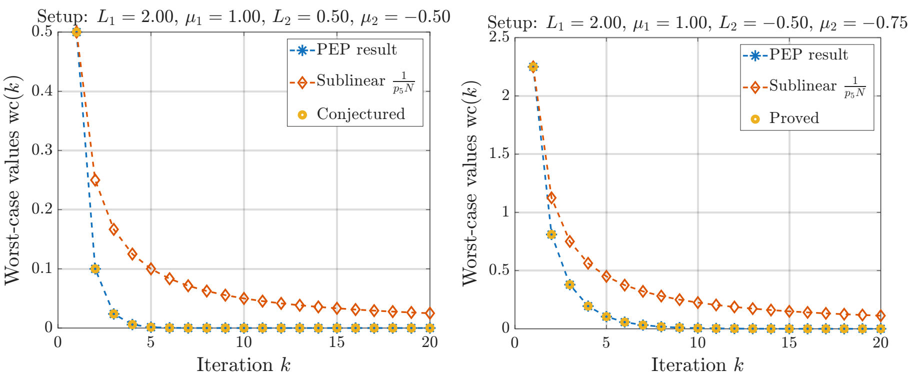

Figure 4.A.3: Examples for regime p_5, with $L_2 > 0$ (*left*) and $L_2 < 0$ (*right*). The *conjectured* part corresponds to the case $L_2 \geq -\mu_2$ (see Conjectures 4.2.3 and 4.2.4), while the *proved* part to $L_2 < -\mu_2$ (see Theorem 4.2.3).

The proof of Theorem 4.2.1 (see Section 4.5.1) provides necessary conditions for the worst-cases corresponding to each of the six regimes. Using these, we elaborate on Proposition 4.2.1 about the tightness of regimes p_1 and p_2, corresponding to the standard DCA, by providing analytical worst-case functions for the cases when both f_1 and f_2 are smooth. The nonsmooth case is discussed in Abbaszadehpeivasti et al. [3, Example 3.1] for regime p_2, with $f_1 \in \mathcal{F}_{0,L_1}$ and $f_2 \in \mathcal{F}_{0,\infty}$.

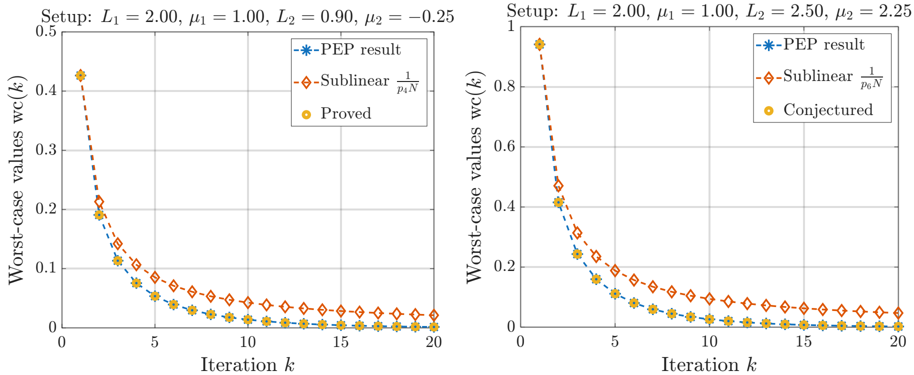

Figure 4.A.4: Examples for regimes p_4 (*left*) and p_6 (*right*), showing that our expressions are non-tight for $N > 1$ iterations.

4.A.1 Worst-case example for regime p_1

Proposition 4.A.1 (Tightness of p_1). *Let f_1^0, f_2^0, g_2^0, x^0 be some real numbers such that $\Delta := f_1^0 - f_2^0 > 0$ and let $N \geq 1$ be an integer. Consider μ_1, L_1, μ_2, L_2 as belonging to the domain of regime p_1 from Table 4.2.1, where the conditions are $0 \leq \mu_1 < L_2 \leq L_1 < \infty$, $\mu_2 < L_1$ and either (i) $\mu_2 \geq 0$ or (ii) $\mu_1 > -\mu_2 > 0$ and $E = \frac{L_2+\mu_2}{L_1 L_2}\frac{L_2-L_1}{-\mu_2} + \mu_1^{-1} - L_1^{-1} \leq 0$. Let $U := -\sqrt{\frac{2\Delta}{p_1 N}}$, with expression of p_1 given in Table 4.2.1. Consider the functions $f_1, f_2 \colon \mathbb{R} \to \mathbb{R}$ defined as:*

$$f_2(x) := \tfrac{1}{2} L_2 (x - x^0)^2 + g_2^0 (x - x^0) + f_2^0$$

$$f_1(x) := \begin{cases} \tfrac{1}{2} L_1 (x - x^0)^2 + g_1^0 (x - x^0) + f_1^0 & x \in (-\infty, x^0]; \\ \tfrac{1}{2} L_1 (x - x^k)^2 + g_1^k (x - x^k) + f_1^k & x \in [x^k, \bar{x}^k]; \\ \tfrac{1}{2} \mu_1 (x - x^{k+1})^2 + g_1^{k+1} (x - x^{k+1}) + f_1^{k+1} & x \in [\bar{x}^k, x^{k+1}]; \\ \tfrac{1}{2} L_1 (x - x^N)^2 + g_1^N (x - x^N) + f_1^N & x \in [x^N, \infty), \end{cases}$$

with $k = 0, 1, \dots, N-1$, $\nabla f_2(x^0) = g_2^0$, $f_1(x^k) = f_1^k$, $\nabla f_1(x^k) = g_1^k$ *and:*

$$x^k = x^0 - k\frac{U}{L_2}, \quad \forall k = 0, \dots, N;$$

$$g_1^k = g_2^0 - (k-1)U, \quad \forall k = 0, \dots, N;$$

$$f_1^k = f_2(x^k) + \frac{N-k}{N}\Delta, \quad \forall k = 0, \dots, N;$$

$$\bar{x}^k = x^k - \frac{L_2 - \mu_1}{L_1 - \mu_1}\frac{U}{L_2}, \quad \forall k = 0, \dots, N-1.$$

Then by performing N iterations of (DCA) *on the function $F(x) = f_1(x) - f_2(x)$, starting from x^0, the following result holds:*

$$\tfrac{1}{2}\min_{0 \le k \le N}\left\{\|\nabla f_1(x^k) - \nabla f_2(x^k)\|^2\right\} = \frac{\Delta}{p_1(L_1, L_2, \mu_1, \mu_2)N}.$$

Proof. By construction, $f_1 \in \mathcal{F}_{\mu_1, L_1}$ and $f_2 \in \mathcal{F}_{\mu_2, L_2}$, where μ_2 can be any real number such that $\mu_2 < L_2$. Moreover, for all $k = 0, 1, \dots, N-1$ we have $g_2^k = g_1^{k+1}$, which is the necessary condition of solving the DCA iteration. Consequently, the quantity $\|g_1^k - g_2^k\|^2$ is exactly $U^2 = \frac{2\Delta}{p_1 N}$, where $\Delta = F(x^0) - F(x^N)$. Given that f_2 is quadratic and f_1 is piecewise quadratic, starting from the initial point x^0, the iterations x^k are uniquely determined by applying (DCA). One can verify that x^k and $\bar{x}^k$ are inflection points for f_1, the curvature changing between μ_1 and L_1. The inflection points $\bar{x}^k$ result from ensuring continuity of both the function and its gradient values and one can verify that $\bar{x}^k \in [x^k, x^{k+1}]$ for all $k = 0, \dots, N-1$. □

The worst-case example from Proposition 4.A.1 builds on the proof of regime p_1 (see Section 4.5.1). Firstly, condition $G = G^+ = L_2\Delta x$ implies $g_1^k - g_2^k = L_2(x^k - x^{k+1})$ at each iteration $k = 0, \dots, N-1$. Additionally, the worst-case after N iterations implies all gradient norms equal with value of $|U|$ from Proposition 4.A.1. Function f_2 is a quadratic of curvature L_2, while f_1 is following a similar construction to the one for gradient descent with stepsizes smaller than $\frac{1}{L}$ in the weakly-convex case (see Proposition 3.6.4).

DCA iterations. The plot on the *right* in Figure 4.A.5 provides an intuitive illustration of the DCA iteration process. Starting from $x^0 = 0$, the iteration moves vertically along the y-axis until intersecting the graph of ∇f_2. Next, it moves horizontally along the x-axis until reaching the graph of ∇f_1 at x^1. This procedure is then repeated for each subsequent iteration.

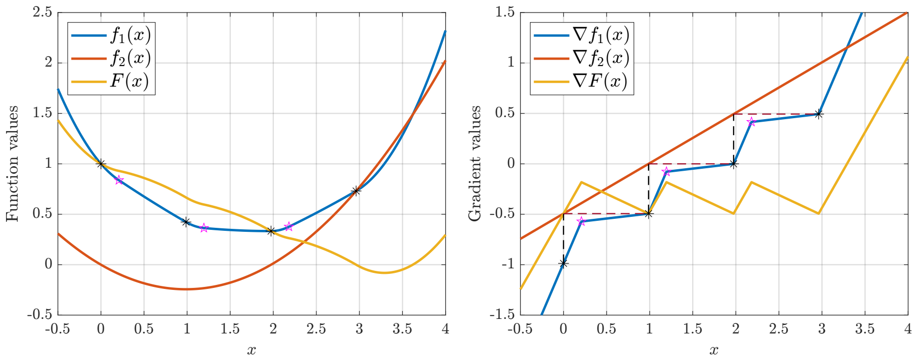

Figure 4.A.5: Worst-case example for regime p_1 after $N = 3$ iterations, showing the function values (*left*) and the gradient values (*right*). Setup: $L_1 = 2$, $\mu_1 = 0.1$, $L_2 = 0.5$, $\mu_2 = -0.01$. The initial condition is $\Delta = f_1(x^0) - f_2(x^0) = 1$. The iterations begin at $x^0 = 0$ where $f_1(x^0) = 1$ and $f_2(x^0) = 0$. The *black* stars represent the DCA iterations x^k, while the *magenta* pentagrams indicate the inflection points of f_1, denoted by $\bar{x}^k$, placed between consecutive iterations x^k.

4.A.2 Worst-case example for regime p_2

Proposition 4.A.2 (Tightness of p_2)**.** *Let f_1^0, f_2^0, g_2^0, x^0 be some real numbers such that $\Delta := f_1^0 - f_2^0 > 0$ and let $N \geq 1$ be an integer. Consider μ_1, L_1, μ_2, L_2 as belonging to the domain of regime p_2 from Table 4.2.1, where the conditions are $\mu_1, \mu_2 \geq 0$ and $\max\{\mu_1, \mu_2\} < L_1 < L_2 < \infty$. Let $U := -\sqrt{\frac{2\Delta}{p_2 N}}$, where the expression of p_2 is given in Table 4.2.1. Consider the functions $f_1, f_2 \colon \mathbb{R} \to \mathbb{R}$ defined as:*

$$f_1(x) := \tfrac{1}{2} L_1 (x - x^0)^2 + g_1^0 (x - x^0) + f_1^0$$

$$f_2(x) := \begin{cases} \tfrac{1}{2}\mu_2 (x - x^0)^2 + g_2^0 (x - x^0) + f_2^0 & x \in (-\infty, x^0]; \\ \tfrac{1}{2}\mu_2 (x - x^k)^2 + g_2^k (x - x^k) + f_2^k & x \in [x^k, \bar{x}^k]; \\ \tfrac{1}{2} L_2 (x - x^{k+1})^2 + g_2^{k+1} (x - x^{k+1}) + f_2^{k+1} & x \in [\bar{x}^k, x^{k+1}]; \\ \tfrac{1}{2}\mu_2 (x - x^N)^2 + g_2^N (x - x^N) + f_2^N & x \in [x^N, \infty), \end{cases} \tag{4.31}$$

where for all $k = 0, 1, \ldots, N-1$ it holds $\nabla f_1(x^0) = g_1^0$, $f_2(x^k) = f_2^k$, $\nabla f_2(x^k) = g_2^k$ and:

$$x^k = x^0 - k\frac{U}{L_1}, \quad \forall k = 0, \ldots, N;$$

$$g_2^k = g_1^0 - (k+1)U, \quad \forall k = 0, \ldots, N;$$

$$f_2^k = f_1(x^k) - \frac{N-k}{N}\Delta, \quad \forall k = 0, \ldots, N;$$

$$\bar{x}^k = x^k - \frac{L_2 - L_1}{L_2 - \mu_2}\frac{U}{L_1}, \quad \forall k = 0, \ldots, N-1.$$

Then by performing N iterations of (DCA) on the function $F(x) = f_1(x) - f_2(x)$, starting from x^0, the following result holds:

$$\tfrac{1}{2} \min_{0 \le k \le N} \{\|\nabla f_1(x^k) - \nabla f_2(x^k)\|^2\} = \frac{\Delta}{p_2(L_1, L_2, \mu_1, \mu_2)N}. \tag{4.32}$$

The proof of Proposition 4.A.2 is similar to the one of Proposition 4.A.1.

Regime p_2 can be understood as a transformation of regime p_1. Specifically, regime p_2 appears when the (DCA) iterations are applied in reverse on the function $-F(x) = f_2(x) - f_1(x)$, with the roles of f_1 and f_2 interchanged. Thus, the dynamic of regime p_1 is mirrored in regime p_2.

4.B Testing curvature shifting (extended results)

In this section we give additional explanatory results corresponding to Section 4.7, which tests the curvature shifting technique.

Recall that for any $x \in \mathbb{R}^n$ we define $f_1(x) = \kappa\|x\|_1 + \eta\frac{\|x\|^2}{2} + \delta_{\bar{B}(0,1)}(x)$ and $f_2(x) \coloneqq \frac{1}{2}x^\top \Sigma x$, where $\kappa > 0$, $\eta \ge 0$, $\delta_{\bar{B}(0,1)}$ is the indicator function over the closed Euclidean unit ball, $\Sigma = A^\top A$ is the sample covariance matrix, where $A \in \mathbb{R}^{20n \times n}$, with $n = 200$ and 10% sparsity. Note that $\mu_1 = \eta$, $L_1 = \infty$, $\mu_2 = \min\{\Lambda(\Sigma)\}$ and $L_2 = \max\{\Lambda(\Sigma)\}$, where $\Lambda(\Sigma)$ denotes the eigenvalues of Σ. Recall $\tilde{f}_i^\lambda = f_i - \lambda\frac{\|\cdot\|^2}{2}$, with $i = \{1, 2\}$, where $\lambda \in (-\infty, \eta]$ is the curvature shift, being the curvature adjusted functions.

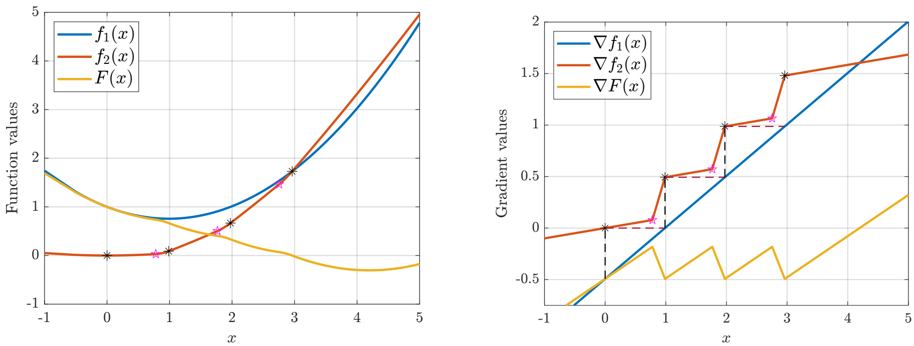

Figure 4.A.6: Worst-case example for regime p_2 after $N = 3$ iterations, showing the function values (*left*) and the gradient values (*right*). Setup: $L_1 = 1.5$, $\mu_1 = 0.25$, $L_2 = 2$, $\mu_2 = 1$. The initial condition is $\Delta = f_1(x^0) - f_2(x^0) = 1$. The iterations begin at $x^0 = 0$ where $f_1(x^0) = 1$ and $f_2(x^0) = 0$. The *black* stars represent the DCA iterations x^k, while the *magenta* pentagrams indicate the inflection points of f_2, denoted by $\bar{x}^k$, placed between consecutive iterations x^k. The dashed lines show the DCA iterations $x^{k+1} = \nabla f_1^*(\nabla f_2(x^k))$.

4.B.1 Closed-form expression of $\partial \tilde{f}_1^{\lambda^*}(y)$

Firstly, for each $y \in \mathbb{R}^n$ we compute $\partial {f_1}^*(y) = \arg\min_{x\in\mathbb{R}^n} \{\langle x, y\rangle - f_1(x)\}$. Let $f(x) := \kappa\|x\|_1 + \delta_{\bar{B}(0,1)}(x)$, such that $f_1 = f + \eta\frac{\|\cdot\|^2}{2}$ and $\tilde{f}_1^\lambda = f + (\eta-\lambda)\frac{\|\cdot\|^2}{2}$. From Themelis et al. [119, §V] we have:

$$\operatorname{prox}_{\gamma f}(y) = \frac{\operatorname{sgn}(y) \odot [|y| - \kappa\gamma\mathbf{1}]_+}{\max\left\{1\,,\, \|\,[|y| - \kappa\gamma\mathbf{1}]_+\,\|_2\right\}},$$

where $\gamma > 0$, $[\cdot]_+$ denotes the positive part, $\odot$ is the elementwise multiplication, and $\mathbf{1}$ is the vector of all ones.

Case $\eta - \lambda > 0$. $\tilde{f}_1^\lambda$ is strongly convex and the subdifferential of the convex conjugate is a singleton. Using the properties of the proximal operator, we get:

$$\nabla \tilde{f}_1^{\lambda^*}(y) = \operatorname{prox}_{(\eta-\lambda)^{-1}f}\left((\eta-\lambda)^{-1}y\right).$$

Case $\eta - \lambda = 0$. We get $\tilde{f}_1^\lambda = f$, which is convex. We obtain

$$\partial f^*(y) = \begin{cases} \mathbf{0}, & \text{if } y = \mathbf{0} \\ \mathbf{0} \cup \mathcal{D}, & \text{if } |y_i| = \kappa,\ \forall i = 1,\ldots,n; \\ \frac{\operatorname{sgn}(y)\odot[|y|-\kappa\mathbf{1}]_+}{\|[|y|-\kappa\mathbf{1}]_+\|_2}, & \text{otherwise}\ , \end{cases}$$

where $\mathcal{D} := \{x \in \bar{B}(0,1) : \operatorname{sgn}(x_i) = \operatorname{sgn}(y_i), \forall i = 1,\dots,n\}$ and $\mathbf{0}$ is the vector of all zeros. Putting together, for all $\lambda \le \eta$ it holds:

$$\partial \tilde{f}_1^{\lambda^*}(y) = \begin{cases} \mathbf{0}, & \text{if } \lambda \le \eta, \quad y = \mathbf{0}; \\ \mathbf{0} \cup \mathcal{D}, & \text{if } \lambda = \eta, \quad |y_i| = \kappa, \forall i = 1,\dots,n; \\ \frac{\operatorname{sgn}(y) \odot [|y| - \kappa \mathbf{1}]_+}{\max\{\eta - \lambda, \|[|y| - \kappa \mathbf{1}]_+\|_2\}}, & \text{otherwise.} \end{cases}$$

The subdifferential is a singleton unless $\lambda = \eta$ and y is on the boundary of the l_∞-ball of radius κ. In this case, in the numerical experiments we select the subgradient $\mathbf{0} \in \partial \tilde{f}_1^{\lambda^*}(y)$.

4.B.2 Extended experimental results

Table 4.B.1: (Case 1) Average number of iterations N_ε to reach ε accuracy of the minimum subgradient norm in the setting $\eta = \mu_1 = 0.5$, $\kappa = 0.02$, corresponding to $\mu_1 > \mu_2$. Lowest values are obtained for $\lambda = \lambda_{\max} = 0.4441$, whereas $\lambda^* = 0.4413$ performs the second best.

λ \ ε	10^{-1}	10^{-2}	10^{-3}	10^{-4}	10^{-5}	10^{-6}	10^{-7}	10^{-8}	10^{-9}	10^{-10}	10^{-11}	10^{-12}
0	1.00	5.74	17.73	46.70	96.70	150.17	242.57	336.27	439.46	544.63	651.33	758.74
$\lambda^* =$ 0.4413	1.00	3.65	8.44	21.49	44.76	70.75	112.57	154.93	200.98	247.93	295.43	343.06
$-0.5\lambda^* =$ -0.2207	1.00	6.83	22.34	59.32	122.73	189.31	307.17	426.39	558.51	693.40	830.17	967.74
$-\lambda^* =$ -0.4413	1.00	8.05	27.00	72.30	147.77	229.13	375.34	521.24	681.09	844.69	1010.64	1177.50
$0.5\lambda^* =$ 0.2207	1.00	4.54	13.06	33.91	70.19	108.43	176.35	244.05	318.63	394.92	472.42	550.15
$\lambda_{\max} =$ 0.4441	1.00	3.63	8.35	21.23	44.01	70.38	111.91	153.78	199.95	246.61	293.76	340.99

Case 2: $\eta < \mu_2$. Parameters: $\eta = 0.2$, $\kappa = 0.02$ (703 runs kept out of 1000). We get $\lambda^* = \mu_1 = 0.2$. Using λ^* to make $\tilde{f}_1^{\lambda}$ convex improves the convergence speed by approximately 20% compared to the initial splitting. Generally, decreasing the curvature of f_2 improves the convergence.

Table 4.B.2: (Case 2) Average number of iterations N_ε to reach ε accuracy of the minimum subgradient norm in the setting $\eta = 0.2$, $\kappa = 0.02$, corresponding to $\eta < \mu_2$. Lowest values are obtained for $\lambda^* = \mu_1 = 0.2$.

λ \ ε	10^{-1}	10^{-2}	10^{-3}	10^{-4}	10^{-5}	10^{-6}	10^{-7}	10^{-8}	10^{-9}	10^{-10}	10^{-11}	10^{-12}
0	1.04	5.88	19.21	49.85	99.16	153.46	245.40	339.67	442.23	547.61	654.99	763.01
$\lambda^* = 0.2$	1.04	4.32	14.23	37.10	75.09	115.90	185.48	256.33	332.81	412.39	492.94	574.18
$-0.5\lambda^* =$ -0.1	1.04	6.48	21.62	55.52	110.98	172.24	276.85	382.11	498.04	616.29	737.17	859.14
$-\lambda^* =$ -0.2	1.04	7.06	23.87	61.44	123.41	192.46	307.80	426.00	554.21	686.04	820.61	956.00
$0.5\lambda^* = 0.1$	1.04	5.26	16.83	43.47	87.38	135.39	215.03	298.14	387.83	479.93	574.02	668.76

4.C PGD invariance to curvature shifting

Proposition 4.C.1 (PGD invariance to curvature shifting)**.** *One PGD iteration with stepsize $\gamma > 0$ on the splitting $F = \varphi + h$ produces the same iterate as one iteration with stepsize $\tilde{\gamma}$ on $F = \tilde{\varphi} + \tilde{h}$, with $\tilde{\varphi} = \varphi + \lambda\frac{\|\cdot\|^2}{2}$, $\tilde{h} = h - \lambda\frac{\|\cdot\|^2}{2}$, $\tilde{\gamma}^{-1} = \gamma^{-1} - \lambda$, where $\lambda < \gamma^{-1}$.*

Proof. By substituting the (subgradients) of the shifted functions

$$\nabla\tilde{\varphi} = \nabla\varphi + \lambda x$$

$$\partial\tilde{h} = \partial h - \lambda x$$

into the PGD iteration

$$\tilde{x}^+ = (\mathrm{id} + \tilde{\gamma}\partial\tilde{h})^{-1}(\mathrm{id} - \tilde{\gamma}\nabla\tilde{\varphi})\tilde{x}$$

we get

$$\tilde{x}^+ = \big[(1-\lambda\tilde{\gamma})\mathrm{id} + \tilde{\gamma}\partial h\big]^{-1}\big[(1-\lambda\tilde{\gamma})\mathrm{id} - \tilde{\gamma}\nabla\varphi\big]\tilde{x},$$

which simplifies to

$$\tilde{x}^+ = \big(\mathrm{id} + \frac{\tilde{\gamma}}{1+\lambda\tilde{\gamma}}\partial h\big)^{-1}\big(\mathrm{id} - \frac{\tilde{\gamma}}{1+\lambda\tilde{\gamma}}\nabla\varphi\big)\tilde{x}.$$

Then the iterates are equivalent for the choice $\gamma = \frac{\tilde{\gamma}}{1+\lambda\tilde{\gamma}}$. □

Corollary 4.C.1 (PPA invariance to curvature shifting)**.** *One proximal point algorithm iteration with stepsize $\gamma > 0$ on the objective $F = h$ produces the same iterate as one PGD iteration with stepsize $\tilde{\gamma}$ on $F = \tilde{\varphi} + \tilde{h}$, with $\tilde{\varphi} = \lambda\frac{\|\cdot\|^2}{2}$, $\tilde{h} = h - \lambda\frac{\|\cdot\|^2}{2}$, $\tilde{\gamma}^{-1} = \gamma^{-1} - \lambda$, where $\lambda < \gamma^{-1}$.*

Proof. The proof results by considering $\varphi = 0$ in Proposition 4.C.1. □

Chapter 5

Douglas–Rachford splitting algorithm in the smooth plus prox-bounded setting

Work not submitted

The results presented in this chapter have not been submitted for publication yet.

5.1 Introduction

This chapter derives exact performance bounds for the Douglas–Rachford splitting (DRS) algorithm applied to composite optimization problems of the form

$$\underset{s\in\mathbb{R}^d}{\text{minimize}}\, \varphi(s) \coloneqq \varphi_1(s) + \varphi_2(s), \tag{5.1}$$

where φ_1 is smooth, i.e., it has a Lipschitz continuous gradient, and φ_2 is l.s.c., regular and prox-bounded (its proximal operator exists for stepsizes below a threshold, see Definition 2.2.3). This framework includes a broad range of problems, convex or nonconvex. The DRS algorithm can be interpreted as a splitting scheme alternating between proximal mappings of φ_1 and φ_2, and gives good performance when each subproblem is typically simple to solve, often amenable to parallel implementations. Our analysis focuses on the case where φ_1 is weakly convex (or hypoconvex).

5.1.1 Literature overview

There is a vast literature in optimization and operator theory on DRS. The method was originally introduced by Douglas and Rachford [40] to solve linear operator equations. Lions and Mercier [78] later established its convergence for monotone inclusions and convex problems. Eckstein and Bertsekas [46] proved that the alternating direction method of multipliers (ADMM) is equivalent to DRS applied to the dual of a linearly constrained convex problem. This insight builds on the foundational work of Rockafellar [97], who established the convergence properties of the Proximal Point Algorithm (PPA), of which DRS is a particular case. Comprehensive discussions of DRS and its applications are available in Bauschke and Combettes [13] and the survey [30]. Giselsson and Boyd [52] show linear convergence conditions for DRS and ADMM under smoothness and strong convexity assumptions, results extended by Moursi and Vandenberghe [82] and tightened by Ryu et al. [104] using performance estimation. Evens et al. [47] present a tight convergence analysis of DRS in the broader class of semimonotone operators, while Abbaszadehpeivasti and Zamani [4] give tight convergence rates for maximally monotone operators.

For convex optimization problems, contraction factors when φ_1 is smooth strongly convex and φ_2 is convex are given in [24, 21] and under more general inexact oracles in [11]. Nguyen et al. [90] provide worst-case rates when φ_1 and φ_2 are both convex, leaving out the smoothness assumption on φ_1.

Nonconvex optimization setups. A key tool is the Douglas–Rachford envelope (DRE), introduced by Patrinos et al. [94], which provides a smooth surrogate of the nonsmooth objective, enabling the study of DRS using gradient-based tools. Li and Pong [76] used this merit function to provide the first convergence analysis in the nonconvex case, in the smooth plus prox-bounded setting, showing empirical improvements over the alternating projection method. Themelis and Patrinos [120] further improved these findings by establishing the DRE properties in nonconvex settings and precisely characterized the parameter regions ensuring monotonic decrease of the DRE, with larger ranges than those of [76]. Their results additionally show a primal equivalence between DRS and ADMM under the same assumptions (one smooth possibly nonconvex and one l.s.c. term). Building on these developments, the current work employs DRE to derive the sharpest progress bounds, in terms of iterate progress or residual subgradient, in the smooth (weakly) convex plus prox-bounded setting.

5.1.2 Contributions

We improve the tight convergence analysis of the DRS method for nonconvex composite problems from [120] and derive explicit upper bounds for (i) the (best) iterate progress and (ii) the (best) residual subgradient measure, valid for any relaxation parameter and constant stepsize guaranteeing monotonic decrease of the DRE. These bounds are conjectured to be tight, supported by extensive Performance Estimation Problem (PEP) experiments that expose one-dimensional worst-case examples. Both metrics exhibit three regimes determined by the curvature bounds of φ_1 and the method parameters: the stepsize and the relaxation.

Inspired by PEP numerical solutions, we identified which interpolation inequalities were sufficient for our analysis. Our constructive proof follows a specific structure and is distinct from the PEP formulation. The key results are the exact decrease bounds after one iteration, given in Theorem 5.3.2 (for iterate progress) and Theorem 5.4.1 (for residual subgradient). Finally, we test the recommendations derived from optimizing the obtained worst-case performance bounds with respect to the stepsize and relaxation parameters. The proofs involve intricate algebraic manipulations and simplifications.

The contributions are summarized in Table 5.1.1.

For ease of verification, we employ a computer algebra system (CAS), namely the MATLAB symbolic toolbox, whose scripts are available in our repository. [1]

Structure. Section 5.2 introduces the problem setup. Section 5.3 and Section 5.4 present the main theorems on exact performance measure bounds, which are proved in Section 5.5. Section 5.6 contains the numerical experiments and Section 5.7 discusses the effect of curvature shifting on DRS.

5.2 Preliminaries

We adopt the notation of [120] and extend their analysis. Algorithm 3 summarizes the iteration scheme; the relevant quantities are defined subsequently.

Assumption 5.2.1 (Assumptions for DRS objective)**.** *The following hold:*

i) $\varphi_1, \varphi_2 \colon \mathbb{R}^d \to \bar{\mathbb{R}}$ *are proper, l.s.c. and extended-real-valued functions.*

[1] Repository: https://github.com/teo2605/Tight_rates_Douglas_Rachford_Splitting

Table 5.1.1: Contributions for Douglas–Rachford splitting in the smooth plus prox-bounded setting. Prior state of the art [120] established tight ranges for convergence parameters (γ, λ) but with non-tight performance constants.

Aspect	**Contributions**
Function class	$\varphi_1 \in \mathcal{F}_{\mu,L}$ with $\mu \le 0$ (allowing $\mu < -L$); φ_2 l.s.c., prox-bounded
Performance metrics	**Iterate progress**: $\mathcal{P}^{\Delta} = \min_k \|\gamma^{-1}(u_k - v_k)\|^2$ **Residual subgradient**: $\mathcal{P}^{\nabla} = \min_k[\mathrm{dist}(0, \hat{\partial}\varphi(v_k))]^2$
Initial condition	$\varphi_\gamma^{\mathrm{DR}}(s_0) - \inf \varphi < \infty$ (finite initial gap in DRE)
One-step bounds	**Iterate progress**: Three regimes $(p_\mu^\Delta, p_L^\Delta, p_3^\Delta)$ in Theorem 5.3.2 **Residual subgradient**: Three regimes $(p_\mu^\nabla, p_L^\nabla, p_3^\nabla)$ in Theorem 5.4.1
N-step rates	$\frac{1}{2}\mathcal{P}^{(\cdot)} \le \frac{\varphi_\gamma^{\mathrm{DR}}(s_0) - \inf \varphi}{\gamma p^{(\cdot)}(\gamma L, \gamma\mu, \lambda)N}$ for $(\cdot) \in \{\Delta, \nabla\}$
Tightness	**Conjectured tight** for all six regimes; extensive PEP validation with one-dimensional worst cases for p_μ and p_L regimes
Optimal tuning	**Iterate progress**: $(\gamma^*, \lambda^*)^\Delta = (\frac{1}{2L}, \frac{3}{2})$ for $\mu \in [-\frac{2L}{7}, 0]$ (Proposition 5.3.1); conjecture for $\mu < -\frac{2L}{7}$ (Conjecture 5.3.1) **Residual subgradient**: Numerical maximization of $\gamma p^\nabla(\gamma L, \gamma\mu, \lambda)$
Improvement	10% (convex, $\mu = 0$) to 46% (smooth, $\mu = -L$) improvement in worst-case constants over [120]
Numerical validation	Theory-recommended tunings rank in top 2–40% on quadratic problems with box/ball constraints
Proof technique	Explicit descent lemmas with algebraic proofs; PEP-informed but analytically complete; computer algebra verification
Curvature shifting	**Non-invariance**: DRS not invariant to curvature shifting (unlike PGD/DCA); counterexample in Section 5.7

ii) The optimization problem has a solution, i.e., $\arg\min \varphi \neq \emptyset$.

iii) $\varphi_1 \in \mathcal{F}_{\mu,L}$ *is* $\max\{-\mu, L\}$*-smooth and* μ*-convex, with* $L > 0$ *and* $\mu \le 0$.

iv) φ *has bounded level sets.*

In comparison to [120], where $|\mu| \le L$, in our work we assume access to curvature bounds, allowing $\mu < -L$, which is a slightly more general assumption. We restrict the tight analysis to $\mu \le 0$. As shown in [120, Remark 3.1],

Algorithm 3: Douglas–Rachford splitting algorithm (DRS)

Data: $\varphi_1 \in \mathcal{F}_{\mu,L}$, with $L > 0$ and $\mu \in (-\infty, 0]$; φ_2 is at least *prox-bounded*; $N \geq 1$ iterations starting from $s_0 \in \mathbb{R}^d$; relaxation $\lambda \in (0, 2)$; stepsize $\gamma \in (0, \min\{\frac{1}{L}, \frac{2-\lambda}{[-2\mu]_-}\})$.

for $k = 0, \ldots, N$ **do**
- Select $u_k \in \mathrm{prox}_{\gamma\varphi_1}(s_k)$;
- Select $v_k \in \mathrm{prox}_{\gamma\varphi_2}(2u_k - s_k)$;
- Take $s_{k+1} = s_k + \lambda\,(v_k - u_k)$

Result: Best iterate (s_k, u_k, v_k) with either
(i) $k = \arg\min_{0\leq k\leq N-1}\{\|\gamma^{-1}(u_k - v_k)\|^2\}$; or
(ii) $k = \arg\min_{0\leq k\leq N-1}\{\left\|\gamma^{-1}(u_k - v_k) - [\nabla\varphi_1(u_k) - \nabla\varphi_1(v_k)]\right\|^2\}$.

Assumption 5.2.1 ensures that φ_2 is prox-bounded with threshold $\gamma_{\varphi_2} \geq \frac{1}{L}$ (recall Definition 2.2.4); the proof requires only the upper curvature bound L.

DRS iteration. Starting from $s_k \in \mathbb{R}^d$, each iteration computes

$$\begin{cases} u_k \in \mathrm{prox}_{\gamma\varphi_1}(s_k); \\ v_k \in \mathrm{prox}_{\gamma\varphi_2}(2u_k - s_k); \\ s_{k+1} = s_k + \lambda\,(v_k - u_k)\,. \end{cases} \tag{DRS}$$

The parameters γ (**stepsize**) and λ (**relaxation**) govern convergence and stability. Themelis and Patrinos [120] derived tight ranges of (γ, λ) for which DRS converges. In particular, they show that the stepsize restriction $\gamma < \frac{1}{L}$ is tight for guaranteeing convergence in the nonconvex setting. Recall that, under Assumption 5.2.1, the prox-boundedness threshold of φ_2 satisfies $\gamma_{\varphi_2} \geq \frac{1}{L}$. Hence, for any stepsize $\gamma < \frac{1}{L}$, we also have $\gamma < \gamma_{\varphi_2}$, which ensures that the $\mathrm{prox}_{\gamma\varphi_2}$ is well defined (i.e., nonempty-valued). Their proof relies on the *Douglas–Rachford envelope* (DRE):

$$\varphi_\gamma^{\mathrm{DR}}(s) := \min_{w\in\mathbb{R}^d}\left\{\varphi_2(w) + \varphi_1(u) + \langle\nabla\varphi_1(u), w - u\rangle + \frac{1}{2\gamma}\|w - u\|^2\right\}, \tag{5.2}$$

where $u = \mathrm{prox}_{\gamma\varphi_1}(s)$. For each s_k, the point v_k is the minimizer of DRE. The DRE is real-valued and strictly continuous for $\gamma < \frac{1}{L}$ [120, Proposition 3.2]. The DRE is closely related to the Forward-Backward Envelope (FBE) [18, 105, 121] through: $\varphi_\gamma^{\mathrm{DR}}(s) = \varphi_\gamma^{\mathrm{FB}}(u)$, where $u = \mathrm{prox}_{\gamma\varphi_1}(s)$.

Several results from [120] are needed in deriving the DRS convergence analysis.

Theorem 5.2.1 (Minimisation equivalence [120, Theorem 3.4])**.** *Under Assumption 5.2.1 and $\gamma < \frac{1}{L}$, the following hold:*

i) $\inf \varphi = \inf \varphi_\gamma^{DR}$*;*

ii) $\arg\min \varphi = \operatorname{prox}_{\gamma\varphi_1}\left(\arg\min \varphi_\gamma^{DR}\right)$*;*

iii) φ *is level bounded if and only if so is* φ_γ^{DR}*.*

Theorem 5.2.2 (Subsequential convergence of DRS [120, Theorem 4.3])**.** *Suppose that Assumption 5.2.1 is satisfied, and consider a sequence* $(s_k, u_k, v_k)_{k\in\mathbb{N}}$ *generated by* (DRS) *with relaxation* $\lambda \in (0, 2)$ *and stepsize* $\gamma \in \min\{\frac{1}{L}, \frac{2-\lambda}{[-2\mu]_-}\}$*, starting from* $s_0 \in \mathbb{R}^d$*. The following hold:*

i) The residual $(u_k - v_k)_{k\in\mathbb{N}}$ *vanishes with rate* $\min_{0\le k\le N-1} \|u_k - v_k\|^2 = o(\frac{1}{k})$*.*

ii) $(u_k)_{k\in\mathbb{N}}$ *and* $(v_k)_{k\in\mathbb{N}}$ *have the same cluster points, all of which are stationary for* φ *and on which* φ *has the same (finite) value, this being the limit of* $\left(\varphi_\gamma^{DR}(s_k)\right)_{k\in\mathbb{N}}$*. In fact, for each* k *one has* $\operatorname{dist}(0, \hat\partial\varphi(v_k)) \le \frac{1-\gamma\mu}{\gamma}\|u_k - v_k\|$*.*

iii) If φ *has bounded level sets, then the sequence* $(s_k, u_k, v_k)_{k\in\mathbb{N}}$ *is bounded.*

Assumption 5.2.1iv) on bounded level sets of φ ensures bounded level sets for $\varphi_\gamma^{\mathrm{DR}}$ (Theorem 5.2.1iii)), implying that the sequence $(s_k, u_k, v_k)_{k\in\mathbb{N}}$ is bounded (Theorem 5.2.2iii)) and thus enabling PEP analysis, which implicitly assumes finite iterates. Theorem 5.2.2 is the base result we extend in the current work.

Performance criteria. The proof of [120, Theorem 4.3] establishes that DRE acts as a Lyapunov function for the admissible stepsize range, reducing the exact performance analysis to quantifying the tightest decrease of the DRE per iteration.

Proposition 5.2.1 (Optimality condition)**.** *Consider the mapping* $F \coloneqq \mathrm{id} - \gamma\nabla\varphi_1$*. Under Assumption 5.2.1, it holds that* $\gamma^{-1}[F(u_k) - F(v_k)] \in \hat\partial\varphi(v_k)$ *and*

$$\operatorname{dist}\left(0, \hat\partial\varphi(v_k)\right) \le \|\gamma^{-1}[F(u_k) - F(v_k)]\| \le (1-\gamma\mu)\|\gamma^{-1}(u_k - v_k)\|.$$

Proof. By definition, $v_k \in \operatorname{prox}_{\gamma\varphi_2}[u_k - \gamma\nabla\varphi_1(u_k)]$ and the optimality condition of the proximal mapping implies

$$\gamma^{-1}(u_k - v_k) - \nabla\varphi_1(u_k) \in \hat\partial\varphi_2(v_k).$$

Adding $\nabla\varphi_1(v_k)$ to both sides, it rewrites as

$$\gamma^{-1}(u_k - v_k) - [\nabla\varphi_1(u_k) - \nabla\varphi_1(v_k)] \in \nabla\varphi_1(v_k) + \hat\partial\varphi_2(v_k) = \hat\partial\varphi(v_k).$$

Therefore, $\gamma^{-1}[F(u_k) - F(v_k)] \in \hat{\partial}\varphi(v_k)$. Since the mapping F is $(1-\gamma\mu)$-Lipschitz continuous, the second inequality follows. □

We improve the convergence rate from Theorem 5.2.2, providing a tight constant for Theorem 5.2.2i), hereafter referred to as *iterate progress* (or *norm*):

$$\mathcal{P}^{\Delta} := \min_{0\leq k\leq N-1}\{\|\gamma^{-1}(u_k - v_k)\|^2\},$$

and a tighter bound for $\left[\operatorname{dist}(0,\hat{\partial}\varphi(v_k))\right]^2$ in Theorem 5.2.2ii), further on called *best subgradient (norm)*:

$$\mathcal{P}^{\nabla} := \min_{0\leq k\leq N-1}\{\|\gamma^{-1}(u_k - v_k) - [\nabla\varphi_1(u_k) - \nabla\varphi_1(v_k)]\|^2\}.$$

We establish, for $(\cdot) \in \{\Delta, \nabla\}$:

$$\frac{1}{2}\mathcal{P}^{(\cdot)} \leq \frac{\varphi_\gamma^{\mathrm{DR}}(s_0) - \inf\varphi}{\gamma p^{(\cdot)}(\gamma L, \gamma\mu, \lambda)N},$$

and conjecture that the functions p^{Δ} and p^{∇} are tight (confirmed numerically via PEP).

5.3 Performance bounds for iterate progress

State of the art. Following [120, Theorem 4.3], a bound on the iterate progress is obtained by telescoping the one-step decrease of the DRE. In our notation, the rate can be written using $p(\gamma L, \gamma\mu, \lambda) = \frac{2\gamma\lambda^2}{(1+\gamma L)^2}c(L,\mu,\gamma,\lambda)$, where c is derived in [120, Theorem 4.1].

Theorem 5.3.1 (State of the art [120, Theorem 4.1])**.** *Let $L > 0$ and $\mu \leq 0$. Consider $\varphi_1 \in \mathcal{F}_{\mu,L}$ and let φ_2 be a proper, l.s.c., and prox-bounded function. Consider one iteration of the DRS algorithm, with relaxation $\lambda \in (0,2)$ and stepsize $\gamma \in (0, \min\{\frac{1}{L}, \frac{2-\lambda}{-2\mu}\}]$, starting from s_0, producing u_0, u_1, v_0, v_1, s_1. Then it holds:*

$$\varphi_\gamma^{\mathrm{DR}}(s_0) - \varphi_\gamma^{\mathrm{DR}}(s_1) \geq \gamma\, p^{\mathit{TP20}}(\gamma L, \gamma\mu, \lambda)\tfrac{1}{2}\left\|\gamma^{-1}(u_0 - v_0)\right\|^2, \tag{5.3}$$

where $p^{\mathit{TP20}} \geq 0$ is defined as

$$p^{\mathit{TP20}}(\ell, m, \lambda) := \begin{cases} \frac{\lambda(2-\lambda+2m)}{(1+\ell)^2} & \text{if } \lambda \in [\frac{2[\ell+m]_+}{\ell}, 2); \\ \frac{\lambda(2-\lambda+2\ell)(1-\ell)}{(1+\ell)^2}; & \text{if } \lambda \in (0, 2(\ell+m)) \\ \frac{\lambda(2-\lambda+\lambda\frac{\ell m}{\ell+m})}{(1+\ell)^2} & \text{if } \lambda \in [2(\ell+m), \frac{2(\ell+m)}{\ell}]. \end{cases} \tag{5.4}$$

5.3.1 Convergence rate for iterate progress

We next provide a sharper one-step bound, proved in Section 5.5.4.

Theorem 5.3.2 (One-step bound for iterate progress). *Let $L > 0$ and $\mu \leq 0$. Consider $\varphi_1 \in \mathcal{F}_{\mu,L}$ and let φ_2 be a proper, l.s.c., and prox-bounded function. Consider one iteration of the DRS algorithm, with relaxation $\lambda \in (0, 2)$ and stepsize $\gamma \in (0, \min\{\frac{1}{L}, \frac{2-\lambda}{-2\mu}\}]$, starting from s_0, producing u_0, u_1, v_0, v_1, s_1. Then it holds:*

$$\varphi_\gamma^{DR}(s_0) - \varphi_\gamma^{DR}(s_1) \geq \gamma\, p^\Delta(\gamma L, \gamma\mu, \lambda) \tfrac{1}{2} \|\gamma^{-1}(u_0 - v_0)\|^2, \tag{5.5}$$

where $p^\Delta \geq 0$ is defined as

$$p^\Delta(\ell, m, \lambda) := \begin{cases} p_\mu^\Delta(\ell, m, \lambda), & \lambda \in [\max\{\frac{4(1+\ell)(1+m)}{4-(1-\ell)(1-m)}, \frac{2(1+m)}{1-m}\}, 2(1+m)]; \\ p_L^\Delta(\ell, m, \lambda), & \lambda \in (0, \min\{\frac{4(1+\ell)(1+m)}{4-(1-\ell)(1-m)}, 1+\ell\}]; \\ p_3^\Delta(\ell, m, \lambda), & \lambda \in [1+\ell, \frac{2(1+m)}{1-m}], \end{cases} \tag{5.6}$$

with

$$p_\mu^\Delta(\ell, m, \lambda) := \frac{\lambda(2 - \lambda + 2m)(1 - m)}{(1 + m)^2}; \tag{p_μ^Δ}$$

$$p_L^\Delta(\ell, m, \lambda) := \frac{\lambda(2 - \lambda + 2\ell)(1 - \ell)}{(1 + \ell)^2}; \tag{p_L^Δ}$$

$$p_3^\Delta(\ell, m, \lambda) := \frac{4 - \lambda^2 - 2(1 + \ell - \lambda)(\lambda + \frac{2(1+\ell-\lambda)}{1-\ell+m(3+\ell)})}{3 + \ell}. \tag{p_3^Δ}$$

Up to the scaling in γ, the leading term in the denominator only depends on γL and $\gamma\mu$, similarly to what we observed earlier for the (proximal) gradient descent method and, implicitly, for the difference-of-convex algorithm. This arises from the smoothness of the mapping $\mathrm{id} - \gamma\nabla\varphi_1$.

Regimes p_μ^Δ and p_L^Δ are associated with worst-case functions of extreme curvature, namely μ and L, respectively. On the other hand, due to the condition $\gamma \in (0, \frac{1}{L}]$, regime p_3^Δ is possible only when $\lambda \in (1, 2)$.

The p_L^Δ branch recovers part of Theorem 5.3.1, but on a larger admissible set. Moreover, regime p_μ^Δ has common terms in the numerator with the first branch from Theorem 5.3.1 up to scaling with $1-\gamma\mu$, which is the Lipschitz constant of mapping F from Proposition 5.2.1, whereas the denominator was approximated by using a conservative bound depending on the Lipschitz constant of φ_1.

Theorem 5.3.3 (N-step rate in iterate progress)**.** *Let $L > 0$ and $\mu \leq 0$. Consider $\varphi_1 \in \mathcal{F}_{\mu,L}$ and let φ_2 be a proper, l.s.c., and prox-bounded function. Consider $N \geq 1$ iterations of the DRS algorithm, with relaxation $\lambda \in (0,2)$ and stepsize $\gamma \in (0, \min\{\frac{1}{L}, \frac{2-\lambda}{-2\mu}\}]$, starting from s_0, generating the sequence $(s_k, u_k, v_k)_{k=\{0,\dots,N+1\}}$. Then, with $p^\Delta(\gamma L, \gamma\mu, \lambda)$ as defined in Theorem 5.3.2, it holds that:*

$$\frac{1}{2}\min_{0\leq k\leq N-1}\{\|\gamma^{-1}(u_k - v_k)\|^2\} \leq \frac{\varphi_\gamma^{DR}(s_0) - \inf\varphi}{\gamma p^\Delta(\gamma L, \gamma\mu, \lambda)N}. \tag{5.7}$$

Proof. Telescoping the inequality (5.5) from Theorem 5.3.2 for N iterations yields (after flipping the order)

$$\gamma p^\Delta(\gamma L, \gamma\mu, \lambda)\frac{1}{2}\sum_{k=0}^{N}\|\gamma^{-1}(u_k - v_k)\|^2 \leq \varphi_\gamma^{\mathrm{DR}}(s_0) - \varphi_\gamma^{\mathrm{DR}}(s_{N+1}).$$

The conclusion results by lower bounding the l.h.s. by taking the minimum norm across iterations and upper bounding the r.h.s. by using $\varphi_\gamma^{\mathrm{DR}}(s_{N+1}) \geq \inf\varphi_\gamma^{\mathrm{DR}} = \inf\varphi$, where the identity follows from Theorem 5.2.1i). ☐

Remark 5.3.1 (On tightness for iterate progress regimes)**.** *PEP evidence indicates simple one-dimensional worst cases for regimes p_μ^Δ and p_L^Δ (but not tight for p_3^Δ) with*

$$\nabla\varphi_1(u_0) - \nabla\varphi_1(u_1) = \begin{cases} \mu\,(u_0 - u_1), & \textit{for regime } p_\mu^\Delta; \\ L\,(u_0 - u_1), & \textit{for regime } p_L^\Delta. \end{cases} \tag{5.8}$$

This extra information is exploited in the proof of Theorem 5.3.2. Leveraging these numerical findings, one may search for piecewise quadratic functions building φ_1, of curvature μ for regime p_μ^Δ and curvature L for regime p_L^Δ, that should connect between them and also with regime p_3^Δ on the boundaries.

Corollary 5.3.1 shows the particular case of $\lambda = 1$, which corresponds to the standard DRS. In this case, regime p_3^Δ vanishes.

Corollary 5.3.1 (Sublinear rates for standard DRS (iterate progress))**.** *Let $L > 0$ and $\mu \leq 0$. Consider $\varphi_1 \in \mathcal{F}_{\mu,L}$ and let φ_2 be a proper, l.s.c., and prox-bounded function. Consider $N \geq 1$ iterations of the DRS algorithm, with* *relaxation* $\lambda = 1$ *and stepsize* $\gamma \in (0, \min\{\frac{1}{L}, \frac{1}{-2\mu}\}]$*, starting from* s_0*, generating the sequence* $(s_k, u_k, v_k)_{k=\{0,\dots,N+1\}}$*. Then it holds that:*

$$\frac{1}{2} \min_{0 \le k \le N-1} \{\|\gamma^{-1}(u_k - v_k)\|^2\} \le \frac{\varphi_\gamma^{DR}(s_0) - \inf \varphi}{\gamma \min\{\frac{(1+2\gamma\mu)(1-\gamma\mu)}{(1+\gamma\mu)^2}, \frac{(1+2\gamma L)(1-\gamma L)}{(1+\gamma L)^2}\} N}. \quad (5.9)$$

Moreover, if $\mu = -L$*, i.e., no lower curvature bound on* φ_1 *is known, then for any* $\gamma \in (0, \frac{1}{2L})$ *it holds*

$$\frac{1}{2} \min_{0 \le k \le N-1} \{\|\gamma^{-1}(u_k - v_k)\|^2\} \le \frac{\varphi_\gamma^{DR}(s_0) - \inf \varphi}{\gamma \min\{\frac{(1-2\gamma L)(1+\gamma L)}{(1-\gamma L)^2}, \frac{(1+2\gamma L)(1-\gamma L)}{(1+\gamma L)^2}\} N}. \quad (5.10)$$

5.3.2 Optimal tuning for iterate progress

Examining Theorem 5.3.3, we assume that the effect of γ on the initial evaluation $\varphi_\gamma^{\mathrm{DR}}(s_0)$ is negligible when optimizing for N iterations, with N large. Then the best tuning maximizes the denominator's leading term $\gamma p^\Delta(\gamma L, \gamma\mu, \lambda)$. The proof of Theorem 5.3.2, given in Section 5.5.4, shows the following interplay between the regimes:

- p_3^Δ is active only if $p_3^\Delta \le \min\{p_\mu^\Delta, p_L^\Delta\}$;
- otherwise, the rate is governed by $\min\{p_\mu^\Delta, p_L^\Delta\}$.

The maximizer lies on regime boundaries, although p^Δ is not exactly $\min\{p_\mu^\Delta, p_L^\Delta, p_3^\Delta\}$.

Proposition 5.3.1 (Optimal tuning for $\mu \in [-2L/7, 0]$ (iterate progress))**.** *Let* $L > 0$ *and* $\mu \in [-\frac{2L}{7}, 0]$*. Consider* $\varphi_1 \in \mathcal{F}_{\mu,L}$ *and let* φ_2 *be a proper, l.s.c., and prox-bounded function. Then the optimized performance bound for the iterate progress after* $N \ge 1$ *iterations of* (DRS) *with stepsize* γ *and relaxation* λ *is achieved at* $(\gamma^*, \lambda^*)^\Delta = (\frac{1}{2L}, \frac{3}{2})$*, yielding the guarantee*

$$\frac{1}{2} \min_{0 \le k \le N-1} \{\|u_k - v_k\|^2\} \le \frac{\varphi_\gamma^{DR}(s_0) - \inf \varphi}{N}.$$

The proof of Proposition 5.3.1 is given in Section 5.5.4.

Conjecture 5.3.1 (Optimal tuning for $\mu < -2L/7$ (iterate progress))**.** *Let* $L > 0$ *and* $\mu \in (-\infty, -\frac{2L}{7})$*. Consider* $\varphi_1 \in \mathcal{F}_{\mu,L}$ *and let* φ_2 *be a proper, l.s.c.,* *and prox-bounded function. Then the smallest value of the guaranteed (scaled) iterate progress after $N \geq 1$ iterations is achieved at the parameters*

$$(\gamma^*, \lambda^*)^\Delta := \left(\frac{h^*}{L}, \frac{4(1+h^*)\left(1+\frac{\mu}{L}h^*\right)}{4-(1-h^*)(1-\left(\frac{\mu}{L}h^*\right)} \right),$$

where $h^* = \arg\max_{h\in(0,1)} \left\{ \frac{8h(1-h^2)[1-\left(\frac{\mu}{L}h\right)^2]}{[4-(1-h)\left(1-\frac{\mu}{L}h\right)]^2} \right\}$.

The motivation behind Conjecture 5.3.1 is that the optimizer lies on the junction between p^Δ_μ and p^Δ_L, where it holds

$$\lambda = \frac{4(1+\gamma L)(1+\gamma\mu)}{4-(1-\gamma L)(1-\gamma\mu)}.$$

With $h = \gamma L$, it remains only to solve the maximization problem defining h^*.

5.3.3 Illustrations of the tight denominator for iterate progress

Figure 5.3.1 depicts the contour plots of $\gamma L p^\Delta(\gamma L, \gamma\mu, \lambda)$ (the denominator amplified by L) from Theorem 5.3.2 for various curvature ratios $\frac{\mu}{L}$, on the entire feasible domain of (normalized) stepsize γL and relaxation parameter λ. The three different regimes are highlighted, together with the boundary conditions delimiting them. We also show the largest value in each contour plot, which follows the optimal tuning results from Proposition 5.3.1 and Conjecture 5.3.1.

Comparison with the state of the art. We depict in Figure 5.3.2 the relative improvement $\frac{p^\Delta - p^{\mathrm{TP20}}}{p^{\mathrm{TP20}}}$, while in Table 5.3.1 we compare the optimized tuning for both rates in terms of (i) best guarantee improvement $\frac{\max \gamma p^\Delta}{\max \gamma p^{\mathrm{TP20}}}$ and (ii) actual best guarantee improvement $\frac{\max \gamma p^\Delta}{\gamma p^\Delta|_{\arg\max \gamma p^{\mathrm{TP20}}}}$. In the first case, we compare the optimized guarantees, whereas in the second one the optimized performance with the best tuning from [120].

We observe a theoretical worst-case improvement of 10% in the convex case ($\mu = 0$), growing as μ decreases. In the smooth case ($\mu = -L$) the improvement is about 46%. There is also a qualitative difference between areas of the best tuning: in our bound, $\lambda^* < 1$ only if $\mu < -1$, while p^{TP20} gives $\lambda^* < 1$ already for mild weak-convexity.

The discontinuous jumps at $\mu = -0.63$ and $\mu = -\frac{2}{3}$ for optimal tuning in p^{TP20} indicate conservatism in the one regime's analysis.

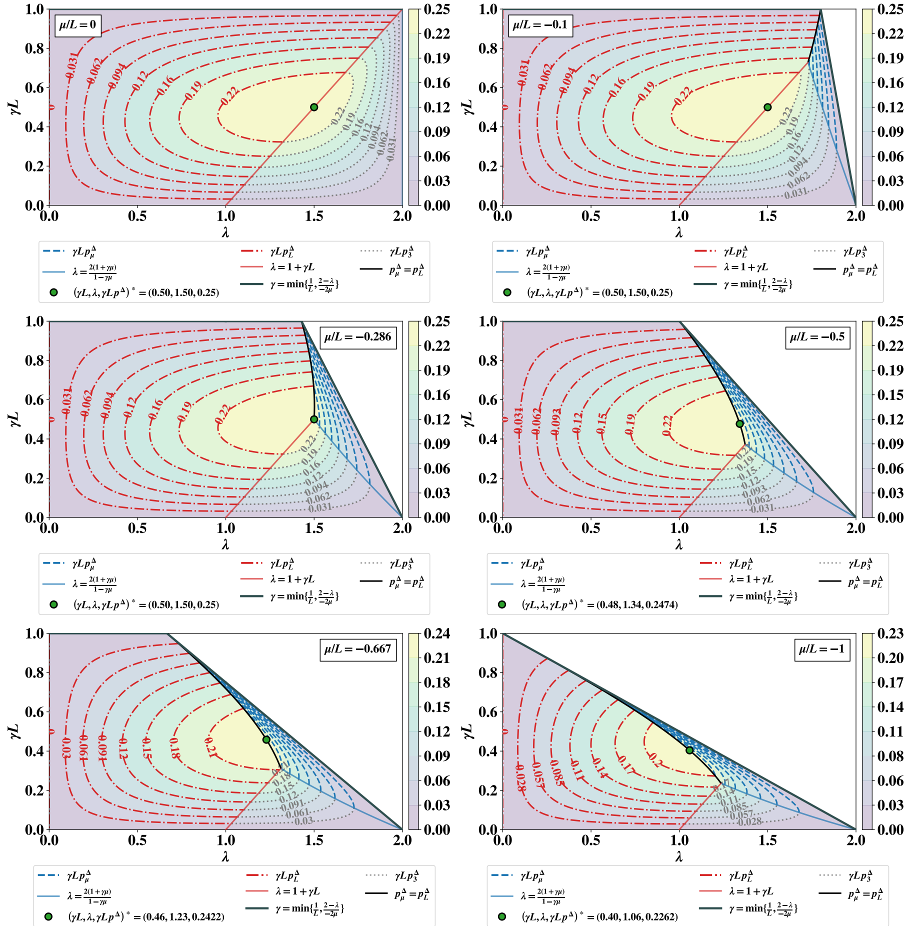

Figure 5.3.1: Contour plots of the tight denominator $\gamma L p^{\Delta}(\gamma L, \gamma\mu, \lambda)$ for various curvature ratios $\frac{\mu}{L}$, on the entire feasible domain. Optimal theoretical tuning is marked by the green dot. The optimal tuning always belongs to the regime p_L^{Δ} and is constant $(\lambda^*, (\gamma L)^*)^{\Delta} = (1.5, 0.5)$, when $\frac{\mu}{L} \in [\frac{-2}{7}, 0]$, or lies on the border with p_μ^{Δ}, otherwise (see Conjecture 5.3.1).

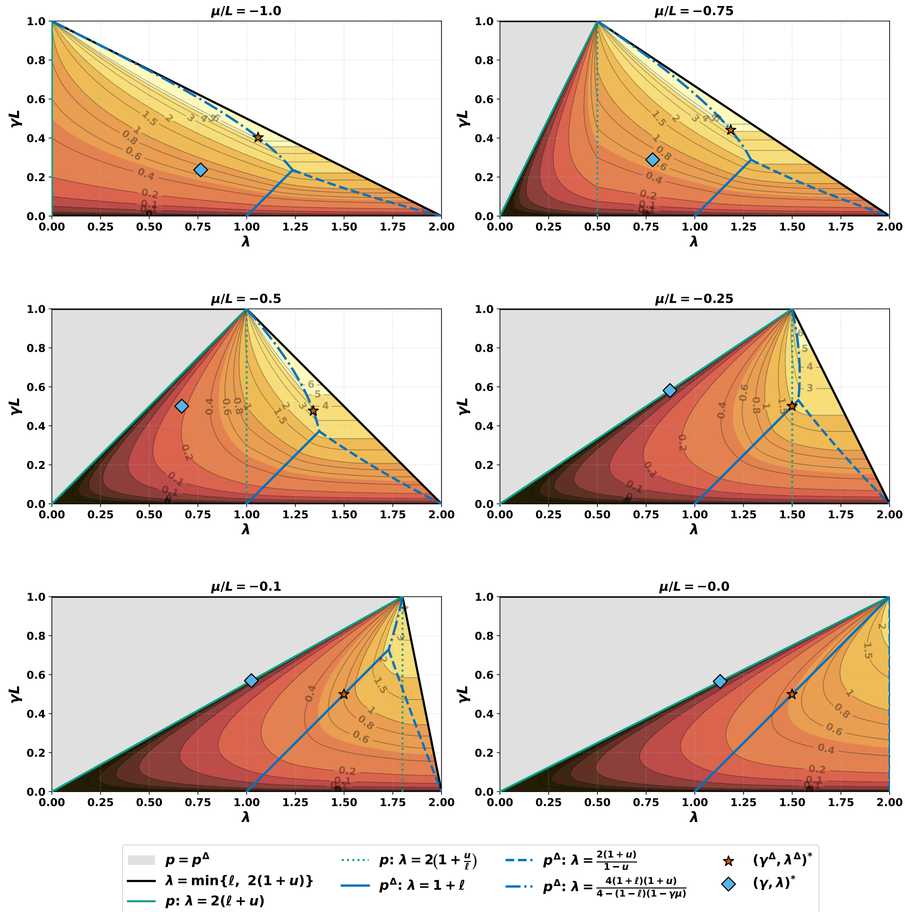

Figure 5.3.2: Relative improvement of our bound over [120], shown as the ratio $\frac{p^{\Delta}-p^{\mathrm{TP20}}}{p^{\mathrm{TP20}}}$ for several curvature ratios μ/L. Gray regions indicate equality of the two performance bounds. Blue curves delineate the regime boundaries of p^{Δ}; green curves delineate those of p^{TP20}. Markers display the optimizing parameter pairs for each method, $((\gamma L)^*, \lambda^*)^{\Delta}$ and $((\gamma L)^*, \lambda^*)^{\mathrm{TP20}}$, respectively.

Table 5.3.1: Improvements at the optimized worst case of our bound relative to the state of the art (here $L = 1$ for simplicity). The quantity $\max \gamma\, p^{\Delta}$ denotes the maximized denominator in our best–iteration–distance bound (see Theorem 5.3.2; the explicit mention of γ, λ is omitted). Similarly, $\max \gamma\, p^{\mathrm{TP20}}$ denotes the corresponding maximized denominator from the state-of-the-art result [120, Theorem 4.1], with maximizer $(\gamma^*, \lambda^*)^{\mathrm{TP20}} \in \arg\max \gamma\, p^{\mathrm{TP20}}$. For completeness, we also report the comparison versus the realized value $\gamma\, p^{\Delta}\big|_{\arg\max \gamma p^{\mathrm{TP20}}}$, i.e., our denominator evaluated at $(\gamma^*, \lambda^*)^{\mathrm{TP20}}$.

μ	$(\gamma^*, \lambda^*)^{\Delta}$	$(\gamma^*, \lambda^*)^{\mathrm{TP20}}$	$\frac{\max \gamma p^{\Delta}}{\max \gamma p^{\mathrm{TP20}}}$	$\frac{\max \gamma p^{\Delta}}{\gamma p^{\Delta}\vert_{\arg\max \gamma p^{\mathrm{TP20}}}}$
0	(0.5, 1.5)	(0.561, 1.123)	1.102	1.102
-0.1	(0.5, 1.5)	(0.570, 1.027)	1.159	1.159
-0.25	(0.5, 1.5)	(0.580, 0.870)	1.285	1.285
-2/7	(0.5, 1.5)	(0.585, 0.836)	1.326	1.325
-1/3	(0.496, 1.463)	(0.588, 0.784)	1.387	1.387
-0.4	(0.490, 1.413)	(0.568, 0.724)	1.487	1.429
-0.5	(0.478, 1.341)	(0.5, 0.666)	1.669	1.430
-0.55	(0.471, 1.307)	(0.467, 0.636)	1.781	1.454
-0.6	(0.466, 1.273)	(0.434, 0.604)	1.913	1.496
-0.63	(0.463, 1.253)	(0.414, 0.586)	2.005	1.526
-2/3	(0.459, 1.229)	(0.311, 0.792)	2.130	1.338
-0.75	(0.447, 1.179)	(0.288, 0.782)	2.237	1.375
-1	(0.409, 1.051)	(0.236, 0.762)	2.508	1.468
-1.5	(0.333, 0.888)	(0.173, 0.740)	2.865	1.596
-2	(0.271, 0.808)	(0.137, 0.724)	3.070	1.669

5.4 Performance bounds for residual subgradient

We derive a one-step upper bound for the residual subgradient (in Theorem 5.4.1), and by telescoping obtain an N-step rate (in Theorem 5.4.2). Unlike for iterate progress, no closed-form optimal worst-case tuning (λ, γ) is available; we therefore report several numerical maximizers for a range of curvature ratios (see Section 5.4.2).

5.4.1 Convergence rate for residual subgradient

Theorem 5.4.1 (One-step bound for residual subgradient)**.** *Let $L > 0$, $\mu \le 0$ and $\lambda \in (0, 2)$. Consider $\varphi_1 \in \mathcal{F}_{\mu,L}$ and let φ_2 be a proper, l.s.c., and prox-bounded function. Let $\gamma \in (0, \min\{\frac{1}{L}, \frac{2-\lambda}{-2\mu}\}]$. Consider one iteration of the DRS algorithm, starting from s_0, producing u_0, u_1, v_0, v_1, s_1. Then it holds:*

$$\varphi_\gamma^{DR}(s_0) - \varphi_\gamma^{DR}(s_1) \ge \gamma p^{\nabla}(\gamma L, \gamma\mu, \lambda)\tfrac{1}{2}\Big[\operatorname{dist}\big(0, \hat{\partial}\varphi(v_0)\big)\Big]^2,$$

where $p^{\nabla} \ge 0$ is defined as

$$p^{\nabla}(\ell, m, \lambda) = \begin{cases} p_\mu^{\nabla}(\ell, m, \lambda), & 1 - \lambda + m \le 0 \textit{ and } p_\mu^{\nabla} \le p_L^{\nabla}; \\ p_L^{\nabla}(\ell, m, \lambda), & \frac{(2-\lambda)^2 - 2(1-\ell)}{(1-\lambda)^2 + 2\ell(1-\lambda) + 1} \ge m \textit{ and } p_L^{\nabla} \le p_\mu^{\nabla}; \\ p_3^{\nabla}(\ell, m, \lambda), & 1 - \lambda + m \ge 0 \textit{ and } \frac{(2-\lambda)^2 - 2(1-\ell)}{(1-\lambda)^2 + 2\ell(1-\lambda) + 1} \le m, \end{cases} \tag{5.11}$$

with

$$p_\mu^{\nabla}(\ell, m, \lambda) := \frac{\lambda\,(2 - \lambda + 2m)\,(1 - m)}{(1 - m^2)^2}; \tag{p_μ^{∇}}$$

$$p_L^{\nabla}(\ell, m, \lambda) := \frac{\lambda\,(2 - \lambda + 2\ell)\,(1 - \ell)}{[1 - \ell m + (1 - \lambda)(\ell - m)]^2}; \tag{p_L^{∇}}$$

$$p_3^{\nabla}(\ell, m, \lambda) := \frac{\lambda[2(1 + \ell) - \lambda] - \frac{4(\ell - m)}{1-m}}{(1 - \ell)(1 + m)^2 - (\ell - m)(2 - \lambda)^2}. \tag{p_3^{∇}}$$

In Section 5.5.5 we provide a methodical demonstration covering all the parameters in the feasible domain, leading to the analytical expressions of all three regimes. Relative to best iterate progress criterion (Theorem 5.3.2), the numerators of regimes p_μ and p_L are the same. However, p_μ^{∇} carries an extra factor of $(1-\gamma\mu)^2$ in the denominator. For p_L, the denominators differ more significantly, with an extra dependence on $\gamma\mu$ for p_L^∇. We also observe the "residual" regime p_3^∇, that can be active only if $\lambda \in (0,1)$ (due to $\gamma\mu \leq 0$) and additionally requires $\gamma L \leq \frac{1}{2}$.

Theorem 5.4.2 (N-step rate in residual subgradient)**.** *Let $L > 0$ and $\mu \leq 0$. Consider $\varphi_1 \in \mathcal{F}_{\mu,L}$ and let φ_2 be a proper, l.s.c., and prox-bounded function. Consider $N \geq 1$ iterations of the DRS algorithm, with relaxation $\lambda \in (0,2)$ and stepsize $\gamma \in (0, \min\{\frac{1}{L}, \frac{2-\lambda}{-2\mu}\}]$, starting from s_0, generating the sequence $(s_k, u_k, v_k)_{k=\{0,\dots,N+1\}}$. Then, with $p^\nabla(\gamma L, \gamma\mu, \lambda)$ as defined in Theorem 5.4.1, it holds that:*

$$\frac{1}{2}\Big[\min_{0\leq k\leq N-1}\{\operatorname{dist}\big(0, \hat{\partial}\varphi(v_k)\big)\}\Big]^2 \leq \frac{\varphi_\gamma^{DR}(s_0) - \inf\varphi}{\gamma p^\nabla(\gamma L, \gamma\mu, \lambda) N}.$$

We omit its proof, because it follows the same arguments as in Theorem 5.3.3.

Remark 5.4.1 (On tightness for residual subgradient regimes)**.** *PEP evidence supports tightness with one-dimensional worst-case instances and the following linear relations for regimes p_μ^∇ and p_L^∇ (not tight for p_3^∇)*

$$\nabla\varphi_1(u_0) - \nabla\varphi_1(u_1) = \begin{cases} \mu\,(u_0 - u_1), & \textit{for regime } p_\mu^\nabla; \\ L\,(u_0 - u_1), & \textit{for regime } p_L^\nabla. \end{cases}$$

$$\nabla\varphi_1(u_1) - \nabla\varphi_1(v_0) = \begin{cases} -(1 - \frac{1+\gamma\mu}{\lambda})\,[\nabla\varphi_1(u_0) - \nabla\varphi_1(u_1)], & \textit{for regime } p_\mu^\nabla; \\ \frac{-\mu}{L}(1 - \frac{1+\gamma L}{\lambda})\,[\nabla\varphi_1(u_0) - \nabla\varphi_1(u_1)], & \textit{for regime } p_L^\nabla. \end{cases}$$

Similarly to the iterate progress criterion, one may search for piecewise quadratic functions building φ_1, of curvature μ for regime p_μ^∇ and curvature L for regime p_L^∇.

In Corollary 5.4.1 we present the particular case of $\lambda = 1$, which corresponds to the standard DRS. In this case, as for the iterate progress, regime p_3^∇ vanishes.

Corollary 5.4.1 (Sublinear rates for standard DRS (residual subgradient))**.** *Let $L > 0$ and $\mu \leq 0$. Consider $\varphi_1 \in \mathcal{F}_{\mu,L}$ and let φ_2 be a proper, l.s.c., and prox-bounded function. Consider $N \geq 1$ iterations of the DRS algorithm, with relaxation $\lambda = 1$ and stepsize $\gamma \in (0, \min\{\frac{1}{L}, \frac{1}{-2\mu}\}]$, starting from s_0, generating the sequence $(s_k, u_k, v_k)_{k=\{0,\dots,N+1\}}$. Then it holds that:*

$$\frac{1}{2}\Big[\min_{0\leq k\leq N-1}\{\operatorname{dist}\big(0, \hat{\partial}\varphi(v_k)\big)\}\Big]^2 \leq \frac{\varphi_\gamma^{DR}(s_0) - \inf\varphi}{\gamma \min\{\frac{(1+2\gamma\mu)(1-\gamma\mu)}{[1-(\gamma\mu)^2]^2}, \frac{(1+2\gamma L)(1-\gamma L)}{(1-\gamma L\gamma\mu)^2}\} N}.$$

Additionally, if $\mu = -L$, i.e., no lower curvature bound on φ_1 is known, then for any $\gamma \in (0, \frac{1}{2L})$ it holds

$$\tfrac{1}{2}\Big[\min_{0\leq k\leq N-1}\{\operatorname{dist}\big(0, \hat{\partial}\varphi(v_k)\big)\}\Big]^2 \leq \frac{\varphi_\gamma^{DR}(s_0) - \inf\varphi}{\gamma\min\{\frac{(1-2\gamma L)(1+\gamma L)}{[1-(\gamma L)^2]^2}, \frac{(1+2\gamma L)(1-\gamma L)}{[1+(\gamma L)^2]^2}\}N}.$$

5.4.2 Illustrations of the tight denominator for residual subgradient

In Figure 5.4.1 we depict the contour plots of $\gamma L p^{\nabla}(\gamma L, \gamma\mu, \lambda)$ (denominator amplified by L) for several curvature ratios $\frac{\mu}{L}$, on the entire feasible domain. The maximizer always belongs to the intersection of regimes p_μ^{∇} and p_L^{∇}, well-separated from the p_3^{∇} region. Consequently, optimal worst-case tuning maximizes $\min\{p_\mu^{\nabla}, p_L^{\nabla}\}$. However, opposite to the iterate progress, we lack a closed-form optimizer.

The comparative bound from [120], given in Theorem 5.2.2ii), is conservative here, having the same leading term as the iterate progress metric up to $1 - \gamma\mu$; consequently, we do not compare against it.

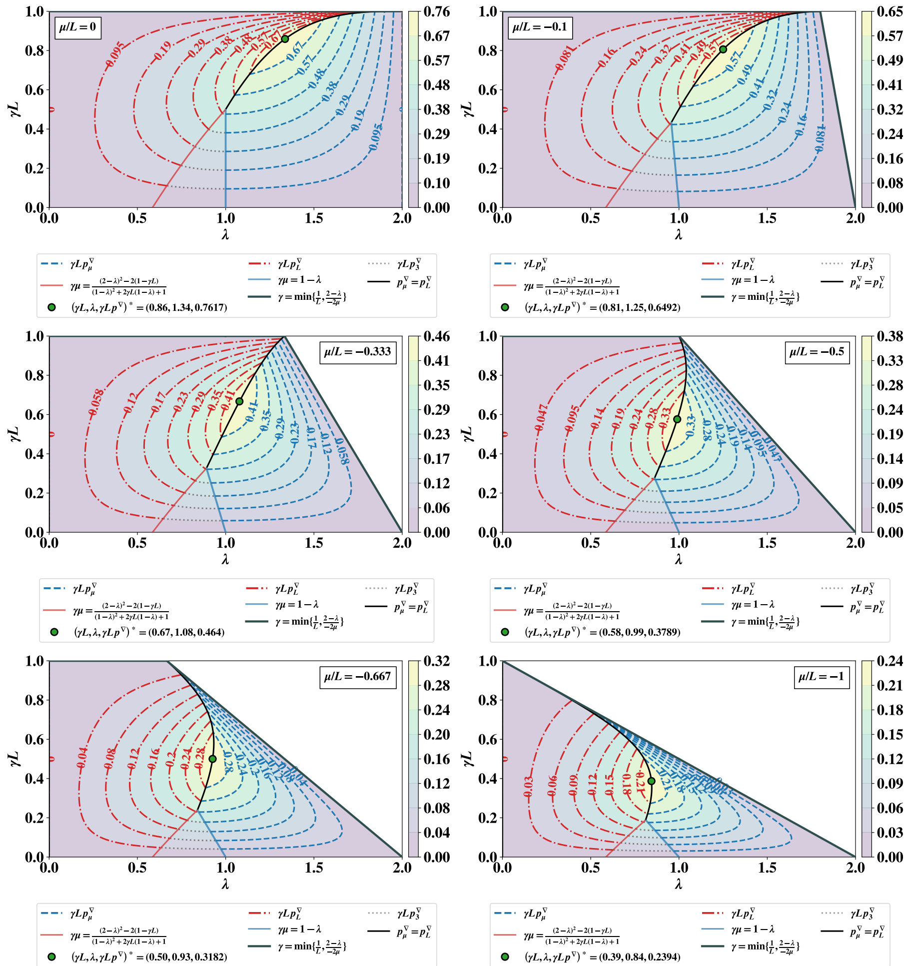

Figure 5.4.1: Contour plots of the tight denominator $\gamma L p^{\nabla}(\gamma L, \gamma\mu, \lambda)$ for the residual subgradient, with various curvature ratios $\frac{\mu}{L}$, highlighting the different active regimes and the boundaries between them. The optimized tuning is depicted as a green dot.

5.5 Proofs

Roadmap of the proofs. The proofs of Theorems 5.3.2 and 5.4.1 follow a common structure. We first introduce the associated PEP Equation (5.12) and discuss the numerical observations. We then establish simplifying notation and key identities in Section 5.5.2. Finally, in Section 5.5.3 we provide a blueprint that reduces the target inequalities to proving positive semidefiniteness of a slack matrix via regime-specific linear transformations; the resulting proofs are carried out in Section 5.5.4 and Section 5.5.5, respectively.

5.5.1 Formulating the PEP

We consider N iterations of (DRS) starting from an initial point s_0. We assume that at the initial point, the value of the DRE is at most at a distance $R > 0$ from the last iteration. As performance metrics we use either $\mathcal{P}^{\Delta} = \min_{0\leq k\leq N-1}\{\|\gamma^{-1}(u_k - v_k)\|^2\}$ or $\mathcal{P}^{\nabla} = \min_{0\leq k\leq N-1}\{\|\gamma^{-1}(u_k - v_k) - [\nabla\varphi_1(u_k) - \nabla\varphi_1(v_k)]\|^2\}$. For fixed parameters N, μ, L, γ and λ, one can solve the PEP (5.12) and obtain a worst-case value for the performance criteria.

$$\begin{aligned} &\underset{s_0,\varphi_1,\varphi_2}{\text{maximize}}\ \mathcal{P}^{\Delta/\nabla}(N, L, \mu, \gamma, \lambda) \\ &\text{subject to} \begin{cases} \text{DRS iteration} \begin{cases} u_k = \operatorname{prox}_{\gamma\varphi_1}(s_k) \\ v_k = \operatorname{prox}_{\gamma\varphi_2}(2u_k - s_k) \\ s_{k+1} = s_k + \lambda(v_k - u_k) \end{cases} \\ \varphi_1 \in \mathcal{F}_{\mu,L} \\ \varphi_2 \text{ prox-bounded (2.2.2)} \\ \varphi^{\mathrm{DR}}(s_0) - \varphi^{\mathrm{DR}}(s_N) \leq R \end{cases} \end{aligned} \tag{5.12}$$

The primal-PEP formulated in (5.12) can be directly plugged into one of the dedicated software packages: PESTO for MATLAB [115] or PEPit for Python [56]. In this way, one can check numerically that our bounds from Theorem 5.3.3 and Theorem 5.4.2 hold. We solved the primal-PEP for many input parameters and observed the worst-case structure mentioned in Remark 5.3.1 and Remark 5.4.1.

The toolboxes also return dual-PEP certificates, which act as numerical proofs. We used these certificates to identify redundant interpolation inequalities that can be pruned without changing the worst-case value. Pruning markedly improves the conditioning of the resulting SDPs: many inequalities are inactive (their multipliers are at machine precision), and in multi-iteration PEPs the accumulated roundoff can otherwise introduce a spurious term. In our experiments, we observed additive term in the denominator, which becomes of the form $pN+q$, with $q \in [10^{-6}, 10^{-3}]$ for $N \geq 10$. After pruning as indicated by the dual solution, q drops to $\leq 10^{-7}$, effectively zero. This supports analyzing a single step and obtaining the tight N-step rate by telescoping, consistent with the DRE serving as a Lyapunov function. The pruning method emphasizes that in each PEP-based analysis one should use problem-dependent procedures and not "blindly" follow naive implementations.

5.5.2 Preliminaries

We focus on one iteration analysis, mapping $(s_0, u_0, v_0) \mapsto (s_1, u_1, v_1)$. We eliminate the variable s by leveraging the DRS iteration: with $k = \{0, 1\}$, $u_k = \operatorname{prox}_{\gamma\varphi_1}(s_k)$ implies $s_k = u_k - \gamma\nabla\varphi_1(u_k)$. We also have $v_k = \operatorname{prox}_{\gamma\varphi_2}(2u_k - s_k) = \operatorname{prox}_{\gamma\varphi_2}(u_k - \gamma\nabla\varphi_1(u_k))$. Then $s_1 = s_0 - \lambda(u_0 - v_0)$ implies

$$u_0 - v_0 = \frac{\gamma}{\lambda}\Big[\gamma^{-1}(u_0 - u_1) + \nabla\varphi_1(u_0) - \nabla\varphi_1(u_1)\Big]. \tag{5.13}$$

Notation. We denote $\ell \coloneqq \gamma L$ and $m \coloneqq \gamma\mu$. We define

$$\Delta u \coloneqq \gamma^{-1}(u_0 - u_1);$$

$$\Delta g_u \coloneqq \nabla\varphi_1(u_0) - \nabla\varphi_1(u_1);$$

$$\Delta g_{u,v} \coloneqq \nabla\varphi_1(u_1) - \nabla\varphi_1(v_0),$$

and all iterates with respect to u_0, $\nabla\varphi_1(u_0)$ and the Δ-quantities, namely $u_1 = u_0 - \gamma\Delta u$, $v_0 = u_0 - \frac{\gamma}{\lambda}(\Delta u + \Delta g_u)$, $\nabla\varphi_1(u_1) = \nabla\varphi_1(u_0) - \Delta g_u$, $\nabla\varphi_1(v_0) = \nabla\varphi_1(u_0) - (\Delta g_{u,v} + \Delta g_u)$.

Consider the set $\mathcal{I}_{u,v} \coloneqq \{(u_0, u_1); (u_1, u_0); (v_0, u_1); (u_0, v_0); (u_1, v_0); (v_0, u_0)\}$, including the pairs of iterates after one iteration; v_1 is not evaluated by φ_1 in the first step. Recall the expression of the DRE at s_k:

$$\varphi_\gamma^{\mathrm{DR}}(s_k) \coloneqq \varphi_2(v_k) + \varphi_1(u_k) + \langle \nabla\varphi_1(u_k), v_k - u_k \rangle + \frac{1}{2\gamma}\|v_k - u_k\|^2.$$

We employ the following quantities relating the iterates and gradients:

$$\begin{aligned} Q_{x,y} \coloneqq\ & \varphi_1(x) - \varphi_1(y) - \langle \nabla\varphi_1(y)\,,\, x - y \rangle - \tfrac{\mu}{2}\|x - y\|^2 \\ & - \tfrac{1}{2(L-\mu)}\|\nabla\varphi_1(x) - \nabla\varphi_1(y) - \mu(x - y)\|^2; \end{aligned}$$

$$PB_{x,y}^{w} \coloneqq \varphi_2(y) - \varphi_2(x) + \frac{1}{2\gamma}\langle y - x\,,\, y + x - 2w \rangle.$$

Since $\varphi_1 \in \mathcal{F}_{\mu,L}$, we have $Q_{x,y} \geq 0$, for all $(x, y) \in \mathcal{I}_{u,v}$ (cf. Theorem 2.2.1). Similarly, since φ_2 is a prox-bounded function, for all pairs $(x, y) \in \mathcal{I}_{u,v}$, with $x = \operatorname{prox}_{\gamma\varphi_2}(w)$, it holds that $PB^{w}_{x,y} \geq 0$ (cf. Theorem 2.2.2).

For brevity of exposition, we treat the vectors of iterations and gradients as scalars when building problem-dependent quadratic forms.

5.5.3 A blueprint of the proofs

We consider the following target inequality, with $p = p(\ell, m, \lambda) \geq 0$:

$$\varphi_\gamma^{\mathrm{DR}}(s_0) - \varphi_\gamma^{\mathrm{DR}}(s_1) - \frac{\gamma p}{2} \|\mathcal{P}\|^2 \\ - \sum_{(\bar{u},\bar{v}) \in \mathcal{I}_{u,v}} \alpha_{(\bar{u},\bar{v})} Q_{(\bar{u},\bar{v})} - \sum_{0 \leq i,j \leq 1} \beta_{i,j} PB^{w_i}_{v_i,v_j} \geq 0, \tag{5.14}$$

where the performance criterion $\mathcal{P}$ is either $\mathcal{P}^{\Delta}$ (smallest iterate progress) or $\mathcal{P}^{\nabla}$ (smallest residual subgradient), that rewrite as:

$$\mathcal{P}^{\Delta} = \|\gamma^{-1}(u_0 - v_0)\|^2 = \|\lambda^{-1}(\Delta u + \Delta g_u)\|^2;$$

$$\begin{aligned} \mathcal{P}^{\nabla} &= \|\gamma^{-1}\left(u_0 - v_0\right) - \left[\nabla\varphi_1(u_0) - \nabla\varphi_1(v_0)\right]\|^2 \\ &= \|\lambda^{-1}(\Delta g_u + \Delta_u) - (\Delta g_{u,v} + \Delta g_u)\|^2, \end{aligned}$$

and $\alpha(\ell, m, \lambda)$ and $\beta(\ell, m, \lambda)$ are nonnegative multipliers. Actually, any nonnegative choice of p, α and β satisfying (5.14) yields a performance bound. The *tightest* one is obtained when the inequality holds with *equality* and corresponds to the maximum possible value of p. It is not necessary to explicitly define all multipliers and it suffices to show their existence. This condition is weaker than for other PEP-based proofs, allowing some flexibility in building up the demonstration.

Dual-PEP observations. Studying the associated multipliers of the interpolation inequalities (the pruning procedure described above), it is sufficient to consider $\beta_{v_1,v_0} = 0$, $\beta_{v_0,v_1} = 1$ and $\alpha_{v_0,u_0} = 0$. Then in both proofs we show that (5.14) reduces to find a positive semidefinite matrix A, depending on p and the multipliers α and β so that

$$\bar{x}^\top \bar{A} \bar{x} \geq 0,$$

where x is either $x = [\Delta u, \Delta g_u]^\top$ (for iterate progress) or $x = [\Delta u, \Delta g_u, \Delta g_{u,v}]^\top$ (for residual subgradient).

Primal-PEP observations. They lead to the worst-case conditions' structure discussed in Remark 5.3.1 and Remark 5.4.1, implying relationships between Δu, Δg_u and, for best subgradient proof, additionally relating $\Delta g_{u,v}$, where certain linear combinations between them vanish. To emphasize these conditions, we apply a linear transformation $\bar{x} = T^{-1}x$ in order to impose these quantities as entries of the x-vector and equivalently, we have to show $x^\top Ax \geq 0$, where $A = T^\top \bar{A}T$. The linear transformation T can be regarded as preconditioner that generates the worst-case scenarios.

Investigating the inequality to prove, some parts are canceled by the worst-case, therefore their corresponding entries in matrix A should vanish. This imposes a system of equations yielding the explicit solution for p and (some of the) multipliers α. Substituting them in matrix A, the problem reduces to solve a small SDP which depends on one of the multipliers.

Finally, for regimes p_3, where a primal solution is not available, the proof exploits the singularity of matrix $\bar{A}$ in the worst-case and maximizes the value of p with respect to a certain multiplier.

We conclude this section by stating an identity involving quantities on the l.h.s. of (5.14), which is used in the proofs for both criteria. Substituting $\beta_{v_1,v_0} = 0$, $\beta_{v_0,v_1} = 1$, and using the relation $v_1 = \operatorname{prox}_{\gamma\varphi_2}(u_1 - \gamma\nabla\varphi_1(u_1))$, yields, after some algebra,

$$\varphi_\gamma^{\mathrm{DR}}(s_0) - \varphi_\gamma^{\mathrm{DR}}(s_1) - PB_{v_0,v_1}^{u_1-\gamma\nabla\varphi_1(u_1)} = \\ \varphi_1(u_0) - \varphi_1(u_1) - \gamma\langle\nabla\varphi_1(u_0)\,,\,\Delta u\rangle + \frac{\gamma}{2}\begin{bmatrix}\Delta u\\ \Delta g_u\end{bmatrix}^\top \begin{bmatrix}\frac{2-\lambda}{\lambda} & 1\\ 1 & \frac{-2}{\lambda}\end{bmatrix}\begin{bmatrix}\Delta u\\ \Delta g_u\end{bmatrix}. \tag{5.15}$$

5.5.4 Proofs for iterate progress criterion

Proof of Theorem 5.3.2. The PEP dual solution reveals that only two interpolation inequalities for φ_1 are active, thus we consider $\alpha := \alpha_{u_1,u_0}(\ell, m, \lambda)$ and $\alpha_{u_0,u_1}(\ell, m, \lambda) = 1 + \alpha$. We search for $\alpha \geq 0$ and p satisfying (5.14), with $\mathcal{P} = \mathcal{P}^\Delta$. Using the identity (5.15), inequality (5.14) becomes

$$\varphi_1(u_0) - \varphi_1(u_1) - \gamma\langle\nabla\varphi_1(u_0)\,,\Delta u\rangle + \frac{\gamma}{2}\begin{bmatrix}\Delta u\\ \Delta g_u\end{bmatrix}^\top \begin{bmatrix}\frac{2-\lambda}{\lambda} - \frac{p}{\lambda^2} & 1 - \frac{p}{\lambda^2}\\ \cdot & \frac{-2}{\lambda} - \frac{p}{\lambda^2}\end{bmatrix}\begin{bmatrix}\Delta u\\ \Delta g_u\end{bmatrix} \\ - (1+\alpha)Q_{u_0,u_1} - \alpha Q_{u_1,u_0} \geq 0. \tag{5.16}$$

Substituting the expressions of Q_{u_0,u_1} and Q_{u_1,u_0} leads to

$$Q_{u_0,u_1} + \alpha(Q_{u_1,u_0} + Q_{u_0,u_1}) = \varphi_1(u_0) - \varphi_1(u_1) - \gamma\langle\nabla\varphi_1(u_0)\,,\Delta u\rangle - \frac{\gamma}{2}\frac{1}{\ell - m}\begin{bmatrix}\Delta u\\ \Delta g_u\end{bmatrix}^\top \begin{bmatrix}\ell m(1+2\alpha) & -\alpha m - (1+\alpha)\ell\\ \cdot & 1+2\alpha\end{bmatrix}\begin{bmatrix}\Delta u\\ \Delta g_u\end{bmatrix}.$$

Replacing in (5.16) and performing the simplifications, the inequality reads as a quadratic form:

$$\frac{\gamma}{2}\begin{bmatrix}\Delta u\\ \Delta g_u\end{bmatrix}^\top \begin{bmatrix}\frac{\lambda(2-\lambda)-p}{\lambda^2} + \frac{\ell m(1+2\alpha)}{\ell-m} & 1 - \frac{p}{\lambda^2} - \frac{\alpha m+(1+\alpha)\ell}{\ell-m}\\ \cdot & -\frac{2\lambda+p}{\lambda^2} + \frac{1+2\alpha}{\ell-m}\end{bmatrix}\begin{bmatrix}\Delta u\\ \Delta g_u\end{bmatrix} \geq 0. \quad (5.17)$$

The PEP observations on the primal problem from Remark 5.3.1 rewrite as

$$\Delta g_u = \begin{cases} m\Delta u, & \text{for regime } p_\mu^\Delta;\\ \ell\Delta u, & \text{for regime } p_L^\Delta.\end{cases}$$

Consequently, we define $t_\mu \coloneqq \Delta g_u - m\Delta u$ and $t_L \coloneqq \Delta g_u - \ell\Delta u$, which cancel on part in the worst-case scenarios: $t_\mu = 0$ for regime p_μ^Δ and $t_L = 0$ for regime p_L^Δ. We change the basis by using the transformation

$$\begin{bmatrix}\Delta u\\ \Delta g_u\end{bmatrix} = \frac{1}{\ell - m}\begin{bmatrix}1 & -1\\ \ell & -m\end{bmatrix}\begin{bmatrix}t_\mu\\ t_L\end{bmatrix}$$

and rewrite the inequality as

$$\frac{\gamma}{2(\ell-m)^2}\begin{bmatrix}t_\mu\\ t_L\end{bmatrix}^\top A\begin{bmatrix}t_\mu\\ t_L\end{bmatrix} \geq 0, \quad (5.18)$$

where A is the matrix of the quadratic form,

$$A \coloneqq \begin{bmatrix}\frac{(1-\ell)[2(1+\ell)-\lambda]}{\lambda} - p\frac{(1+\ell)^2}{\lambda^2} & \alpha(\ell-m) + 1 - m - \frac{2(1-\ell m)}{\lambda} + p\frac{(1+m)(1+\ell)}{\lambda^2}\\ \cdot & \frac{(1-m)[2(1+m)-\lambda]}{\lambda} - p\frac{(1+m)^2}{\lambda^2}\end{bmatrix}.$$

The worst-case which imposes strict equality, hence the following cancellations occur:

1. Case p_μ^Δ: $A_{(1,2)} = 0$ and $A_{(2,2)} = 0$ (factors of t_L are canceled);
2. Case p_L^Δ: $A_{(1,2)} = 0$ and $A_{(1,1)} = 0$ (factors of t_μ are canceled).

Canceling $A_{(2,2)}$ leads to $p = p_\mu^\Delta$, while zeroing $A_{(1,1)}$ to $p = p_L^\Delta$. We have that $p_\mu^\Delta \geq 0$ if $\lambda \leq 2(1+m)$, whereas $p_L^\Delta \geq 0$ is satisfied implicitly by the assumption

$\ell \leq 1$. Replacing these values and using that $A_{(1,2)} = 0$, the inequality to show rewrites as follows in a compact form:

$$\frac{\gamma(1+m)^2(1+\ell)^2}{2\lambda^2(\ell-m)^2}(p_\mu^\Delta - p_L^\Delta)\Big[\delta(p_\mu^\Delta \geq p_L^\Delta)\frac{\|t_L\|^2}{(1+\ell)^2} - \delta(p_\mu^\Delta \leq p_L^\Delta)\frac{\|t_\mu\|^2}{(1+m)^2}\Big] \geq 0,$$

where $\delta(x) = 0$ if x is false and $\delta(x) = 1$ if x is true. Therefore, p_μ^Δ and p_L^Δ are complementary to each other. Additionally, they share the cancellation of $A_{(1,2)}$, which yields

$$\alpha(p) = \frac{1}{\ell-m}\Big[\frac{2(1-\ell m)}{\lambda} - (1-m) - p\frac{(1+\ell)(1+m)}{\lambda^2}\Big].$$

Specializing for the two regimes, we get

$$\alpha_\mu(\ell,m,\lambda) \coloneqq \frac{2}{1+m} - \frac{2}{\lambda} - 1 = \frac{1-m}{1+m} - \frac{2}{\lambda};$$

$$\alpha_L(\ell,m,\lambda) \coloneqq \frac{2}{\lambda} - \frac{2}{1+\ell}.$$

Domains of p_μ^Δ and p_L^Δ. Since α must be nonnegative, the proof of p_μ^Δ holds only if $\lambda \geq 2\frac{1+m}{1-m}$, whereas the proof of p_L^Δ holds only if $\lambda \leq 1+\ell$, both conditions imposing regimes' bounds in the branches from (5.6). Moreover, one can check that the difference $p_\mu^\Delta - p_L^\Delta$ can be expressed as

$$p_\mu^\Delta - p_L^\Delta = (\bar\lambda - \lambda)\frac{\lambda(\ell-m)[4-(1-\ell)(1-m)]}{(1+\ell)^2(1+m)^2},$$

where $\bar\lambda \coloneqq \frac{(1+\ell)(1+m)}{4-(1-\ell)(1-m)}$. We have that $4-(1-\ell)(1-m) \geq 4-(2-\frac{\lambda}{2}) = 2+\frac{\lambda}{2}$, where we use that $\ell \in (0,1]$ and $\lambda \leq 2(1+m)$. Then $p_\mu^\Delta \geq p_L^\Delta$ for $\lambda \leq \bar\lambda$ and $p_\mu^\Delta \geq p_L^\Delta$ for $\lambda \leq \bar\lambda$, hence $\bar\lambda$ exactly delimits the two regimes as shown in their definition (5.6).

It remains to derive the expression of p_3^Δ, corresponding to the domain with $\lambda \in (1+\ell, \frac{2(1+m)}{1-m})$, which can be nonempty only when $\lambda > 1$. In this third scenario, a handy worst-case is not available. Inequality (5.17) holds if the matrix (let us call it $\bar A$) is positive semidefinite and the extreme case is obtained if the matrix is singular; in this case, the PSD condition reduces to check the principal minors. With some algebra (using symbolic toolboxes), from the zero determinant condition we obtain $p(\alpha) = \frac{p_N(\alpha)}{p_D(\alpha)}$, where

$$\begin{aligned} p_N(\alpha) \coloneqq\ & -[\alpha^2(\ell-m) + (1+2\alpha)(1-m)]\lambda^2 + \\ & + 2[2\alpha(1-\ell m) + (1-m)(1+\ell)]\lambda - 4(\ell-m); \\ p_D(\alpha) \coloneqq\ & 2\alpha(1+m)(1+\ell) + 4m + (1-\ell)(1-m). \end{aligned}$$

Further on, we derive $\alpha \geq 0$ as the maximizer of $p(\alpha)$. Its gradient remarkably simplifies to

$$\nabla p(\alpha) = \frac{-2\lambda^2(1+\ell)(1+m)(\ell-m)}{[2\alpha(1+\ell)(1+m)+4m+(1-\ell)(1-m)]^2}(\alpha-\alpha_\mu)(\alpha-\alpha_L).$$

Observe that the critical points are exactly α_μ and α_L and the sign of $\nabla p(\alpha)$ is strictly given by the value of α with respect to α_μ and α_L. Since neither of regimes p_μ^Δ or p_L^Δ is active, we have that both $\alpha_\mu < 0$ and $\alpha_L < 0$, thus $p(\alpha)$ strictly decreases with α. Consequently, its maximum value is obtained at $\alpha = 0$, leading after substitution to $p = p_3^\Delta$. Furthermore, we get

$$\bar{A}_{(1,1)} = \frac{4(\ell-m)(1+\ell-\lambda)^2}{\lambda^2\psi(m,\ell)};$$

$$\bar{A}_{(2,2)} = \frac{(\ell-m)[2(1+m)-\lambda(1-m)]^2}{\lambda^2\psi(m,\ell)},$$

where $\psi(\ell,m) \coloneqq 4m+(1-m)(1-\ell)$ dictates the sign of the principal minors. Using that on this regime we have $\ell \leq \lambda-1$, we get $\psi(\ell,m) \geq \psi(\lambda-1,m) = (2+\lambda)m+(2-\lambda)$, which is positive from the condition $-m \leq \frac{2-\lambda}{2+\lambda}$. Finally, we have the identities

$$\bar{A}_{(1,1)} = (p_L^\Delta - p_3^\Delta)\frac{(1+\ell)^2}{\lambda^2};$$

$$\bar{A}_{(2,2)} = (p_\mu^\Delta - p_3^\Delta)\frac{(1+m)^2}{\lambda^2},$$

highlighting the zero-slack at the points of changing regimes and showing that regime p_3^Δ is complementing exactly the other two cases, in this way ensuring the complete covering of the feasible domain. ☐

The proof is constructive, with clear motivations on how to choose the multipliers based on PEP-informed primal and dual solutions. Nevertheless, regime p_3^Δ shows the maximum guarantee that can be obtained using only the interpolation inequality $Q_{[u_0,u_1]}$.

Remark 5.5.1 (Alternative proof for iterate progress regimes)**.** *An alternative way to prove the regimes p_μ^Δ and p_L^Δ, without employing PEP-based information about the worst-case, is by leveraging the observation leading to regime p_3^Δ, namely that their corresponding multiplier is one of the two critical points of $p(\alpha(\ell,m,\lambda))$.*

Proof of Proposition 5.3.1. Without loss of generality, we set $L = 1$ and denote $h = \gamma L$ (the normalized stepsize) and $\kappa \coloneqq \frac{\mu}{L}$ (the curvature ratio).

Then, for a given κ we aim to maximize $hp^{\Delta}(h, \kappa h, \lambda)$, with $\lambda \in (0, 2)$, $h \leq 1$, $\kappa \in [\frac{-2}{7}, 0]$ and $2(1 + \kappa) \geq \lambda$.

Maximizing the value of p_L^{Δ} in the variable λ, we have

$$\max_{\lambda \in \operatorname{dom} p_L^{\Delta}} hp_L^{\Delta} = \max_{\lambda \in \operatorname{dom} p_L^{\Delta}} \frac{h^2[2(1+h) - \lambda](1-h)}{(1+h)^2}.$$

The objective is a concave quadratic in λ, maximized at $\lambda = 1 + h$, which is exactly the boundary between p_L^{Δ} and p_3^{Δ} (maximum possible λ for p_L^{Δ} and minimum possible for p_3^{Δ}). It follows that

$$\max_{\lambda \in \operatorname{dom} p_L^{\Delta}} hp_L^{\Delta} = h(1-h) \leq \frac{1}{4},$$

with equality only at $h^* = \frac{1}{2}$. Then it implies $\lambda^* = 1 + h^* = \frac{3}{2}$ and the maximized performance guarantee $h^* p_L^{\Delta}(h^*, \kappa h^*, \lambda^*) = \frac{1}{4}$.

For regime p_3^{Δ} we have

$$\frac{dp_3^{\Delta}}{d\lambda} = \frac{2(1-\kappa h)(1+h-\lambda)}{\kappa h(3+h) + (1-h)}.$$

To analyse its sign, we recall that regime p_3^{Δ} corresponds to $\lambda \in [1 + h, \frac{2(1+\kappa h)}{1-\kappa h}]$. The non-empty interval condition is equivalent to $\kappa h(3 + h) + (1 - h) \geq 0$, which is exactly the derivative's denominator. On the other hand, the numerator is negative because of the lower bound of λ. Then $\frac{dp_3^{\Delta}}{d\lambda} \leq 0$ and its maximum is attained at the junction point with p_L^{Δ}, namely $\lambda = 1 + h$, with $\max_{\lambda \in \operatorname{dom} p_3^{\Delta}} hp_3^{\Delta} = h(1 - h)$.

Moreover, replacing $\lambda = 1 + h^* = \frac{3}{2}$ in the non-empty interval condition for p_3^{Δ} leads to $\kappa \geq \frac{-2}{7}$. When the inequality is strict, the optimizer cannot belong to the domain of regime p_μ^{Δ} because it is the maximizer of both other regimes, implying $\lambda < \frac{2(1+\kappa h)}{1-\kappa h}$ (outside of the domain).

For $\kappa = \frac{-2}{7}$ and $h = \frac{1}{2}$, the p_3^{Δ}-interval of λ becomes the singleton $\frac{3}{2}$ and, additionally, regime p_μ^{Δ} becomes active and thus we have the triple intersection of the regimes.

Finally, the guaranteed bound results by substituting $\gamma = \frac{1}{2L}$. □

5.5.5 Proof for best subgradient criterion

Proof of Theorem 5.4.1. Let $\alpha_{(\bar{u},\bar{v})} = \alpha_{(\bar{u},\bar{v})}(\ell, m, \lambda)$, with $(\bar{u}, \bar{v}) \in \mathcal{I}_{u,v}$. We aim to find nonnegative values of these multipliers and of $p \geq 0$ satisfying (5.14), with $\mathcal{P} = \mathcal{P}^{\nabla}$. With $x \coloneqq \big[\Delta u,\, \Delta g_u,\, \Delta g_{u,v}\big]^{\top}$, replacing identity (5.15) in inequality (5.14) yields

$$\varphi_1(u_0) - \varphi_1(u_1) - \gamma\langle\nabla\varphi_1(u_0), \Delta u\rangle + \frac{\gamma}{2}x^{\top}A_{\mathcal{P}}x - \sum_{(\bar u,\bar v)\in\mathcal{I}_{u,v}} \alpha_{(\bar u,\bar v)}Q_{(\bar u,\bar v)} \geq 0, \tag{5.19}$$

where

$$A_{\mathcal{P}} \coloneqq \begin{bmatrix} \frac{2-\lambda}{\lambda} - \frac{p}{\lambda^2} & 1 - \frac{p(1-\lambda)}{\lambda^2} & \frac{p}{\lambda} \\ \cdot & \frac{-2}{\lambda} - \frac{p(1-\lambda)^2}{\lambda^2} & \frac{p(1-\lambda)}{\lambda^2} \\ \cdot & \cdot & -p \end{bmatrix}.$$

The linear combination $\sum_{(\bar u,\bar v)\in\mathcal{I}_{u,v}} \alpha_{(\bar u,\bar v)}Q_{(\bar u,\bar v)}$ must include as function values only $\varphi_1(u_0) - \varphi_1(u_1)$. This condition implies

$$\alpha_{u_0,v_0} = 1 + \alpha_{u_1,u_0} - \alpha_{u_0,u_1};$$

$$\alpha_{u_1,v_0} = \alpha_{u_0,u_1} - \alpha_{u_1,u_0} + \alpha_{v_0,u_1} - 1.$$

Substituting in the linear combination yields

$$\sum_{(\bar u,\bar v)\in\mathcal{I}_{u,v}} \alpha_{(\bar u,\bar v)}Q_{(\bar u,\bar v)} = \varphi_1(u_0) - \varphi_1(u_1) - \gamma\langle\nabla\varphi_1(u_0), \Delta u\rangle - \frac{\gamma}{2(\ell-m)}x^{\top}A_Qx,$$

where the entries in A_Q are

$$A_Q^{(1,1)} = \ell m\Big[(1+2\alpha_{u_1,u_0}) + \frac{2(1-\lambda)}{\lambda^2}\big[(1-\lambda)\alpha_{v_0,u_1} + \lambda(1-\alpha_{u_0,u_1}+\alpha_{u_1,u_0})\big]\Big];$$

$$\begin{aligned} A_Q^{(1,2)} = A_Q^{(2,1)} = &-(\ell-m) - \ell\alpha_{u_1,u_0} - m\alpha_{u_0,u_1} \\ &- m(1-\ell)\frac{1+\alpha_{u_1,u_0}-\alpha_{u_0,u_1}}{\lambda} + \frac{2(1-\lambda)}{\lambda^2}\ell m\alpha_{v_0,u_1}; \end{aligned}$$

$$A_Q^{(1,3)} = A_Q^{(3,1)} = -\ell(1-\alpha_{u_0,u_1}+\alpha_{u_1,u_0}) - \frac{1-\lambda}{\lambda}(\ell+m)\alpha_{v_0,u_1};$$

$$A_Q^{(2,2)} = 1 + 2\alpha_{u_1,u_0} - \frac{2m}{\lambda}(1-\alpha_{u_0,u_1}+\alpha_{u_1,u_0} - \frac{\ell}{\lambda}\alpha_{v_0,u_1});$$

$$A_Q^{(2,3)} = A_Q^{(3,2)} = 1 - \alpha_{u_0,u_1} + \alpha_{u_1,u_0} - \frac{\ell+m}{\lambda}\alpha_{v_0,u_1};$$

$$A_Q^{(3,3)} = 2\alpha_{v_0,u_1}.$$

The linear combination includes, with the opposed coefficient, the inner product $\langle\nabla\varphi_1(u_0), \Delta u\rangle$, which appears in (5.19). Thus, the residual reduces to a quadratic form in Δu, Δg_u and $\Delta g_{u,v}$. We define the slack matrix (we omit writing explicitly its dependence on p and multipliers α)

$$A \coloneqq A_Q + (\ell - m)A_{\mathcal{P}}. \tag{5.20}$$

After substituting everything in (5.19), it remains to find p and the multipliers α such that

$$\frac{\gamma}{2(\ell - m)} x^\top A x \geq 0. \tag{5.21}$$

Further on, the analysis is split based on the PEP-informed worst-case relationships for regimes p_μ^∇ and p_L^∇, as mentioned in Remark 5.4.1. When both proofs break, a third one, of regime p_3^∇, emerges. An important quantity leading the proof is $S \coloneqq p_\mu^\nabla - p_L^\nabla$, with $S \geq 0$ if regime p_L^∇ is active and $S \leq 0$ otherwise.

Case 1. Regime p_μ^∇: The PEP worst-case shows that

$$t_\mu \coloneqq \Delta g_u - m\Delta u = 0;$$

$$r_\mu \coloneqq \Delta g_{u,v} + (1 - \frac{1+m}{\lambda})\Delta g_u = 0.$$

Inspired by these identities, we define the transformation

$$T_\mu \coloneqq \begin{bmatrix} 1 & 0 & 0 \\ m & 1 & 0 \\ sm & s & 1 \end{bmatrix},$$

where $s = \frac{1+m-\lambda}{\lambda}$, such that $x = T_\mu x_\mu^\nabla$, with $x_\mu^\nabla = \left[\Delta u,\, t_\mu,\, r_\mu\right]^\top$. Inequality (5.21) is then equivalent to $\tilde{x}^\top (T_\mu^\top A T_\mu)\tilde{x} \geq 0$. Because the worst-case implies $t_\mu = 0$ and $r_\mu = 0$, a direct recipe to obtain the solution such that the inequality holds exactly is to (i) cancel the contributions of Δu and (ii) choose the left-over parameters such that $A_\mu \coloneqq T_\mu^\top A T_\mu \succeq 0$. This implies $A_\mu^{(1,1)} = (A_\mu)_{(1,2)} = (A_\mu)_{(1,3)} = 0$, leading to the system of equations:

$$\begin{cases} 0 = A_{(1,1)} + 2A_{(1,2)} + 2smA_{(1,3)} + m^2 A_{(2,2)} + 2sm^2 A_{(2,3)} + s^2 m^2 A_{(3,3)}; \\ 0 = A_{(1,2)} + sA_{(1,3)} + mA_{(2,2)} + 2smA_{(2,3)} + s^2 m A_{(3,3)}; \\ 0 = A_{(1,3)} + mA_{(2,3)} + smA_{(3,3)}. \end{cases}$$

After substituting the expressions for the entries of matrix A and performing the simplifications, the system of equations reduces to:

$$\begin{cases} 0 = p(1-m)^2(1+m)^2 - 2\lambda(1-m^2) + \lambda^2(1-m); \\ 0 = \lambda^2(\alpha_{u_0,u_1} + \alpha_{v_0,u_1} - 1) - \lambda[\alpha_{u_0,u_1} - \alpha_{u_1,u_0} + (2+m)\alpha_{v_0,u_1} - (1+2m)] \\ \qquad + (1+m)\alpha_{v_0,u_1} - pm(1-m^2); \\ 0 = \lambda(\alpha_{u_0,u_1} - \alpha_{u_1,u_0} + \alpha_{v_0,u_1} - 1) + (1+m)[p(1-m) - \alpha_{v_0,u_1}]. \end{cases}$$

The solution $p = p_\mu^\nabla$ results directly from the first equation. Fixing $\alpha \coloneqq \alpha_{u_0,u_1}$, we obtain the following parametrized solutions for α_{u_1,u_0} and α_{v_0,u_1}:

$$\alpha_{u_1,u_0} = \frac{(2-\lambda)[\lambda-(1+m)]}{\lambda(1+m)};$$

$$\alpha_{v_0,u_1} = \frac{2}{1+m} - \alpha\frac{\lambda}{\lambda-(1+m)}.$$

Note that $\alpha_{u_1,u_0} \geq 0$ implies $\lambda \geq 1+m$, which is the boundary condition between p_μ^∇ and p_3^∇. Moreover, $p = p_\mu^\nabla \geq 0$ only if $\lambda \leq 2(1+m)$ (the condition from hypothesis), which additional implies that $1+m > 0$. Replacing the above multipliers' expressions in the identities for α_{u_0,v_0} and α_{u_1,v_0}, one gets:

$$\alpha_{u_0,v_0} = \frac{2-\lambda}{1+m} - \frac{2(1-\lambda)}{\lambda} - \alpha;$$

$$\alpha_{u_1,v_0} = \frac{\lambda}{1+m} + \frac{2(1-\lambda)}{\lambda} - \alpha\frac{1+m}{\lambda-(1+m)}.$$

All these multipliers must be nonnegative, therefore

$$\alpha \in \left[0, \min\left\{ \underbrace{\frac{2[\lambda-(1+m)]}{\lambda(1+m)}}_{\alpha_1}; \underbrace{\frac{2-\lambda}{1+m} - \frac{2(1-\lambda)}{\lambda}}_{\alpha_2}; \underbrace{\frac{[\lambda-(1+m)]\left[\frac{\lambda}{1+m} + \frac{2(1-\lambda)}{\lambda}\right]}{1+m}}_{\alpha_3} \right\}\right].$$

One can check the identities $\alpha_3 = \alpha_1 + \frac{(1+m-\lambda)[2(1+m)-\lambda]}{(1+m)^2}$ and $\alpha_2 = \alpha_1 + \frac{2(1+m)-\lambda}{1+m}$. The condition $\lambda \in [1+m, 2(1+m)] \cap (0,2)$ implies that the minimum between the three terms is α_1, hence $\alpha \in [0, 2\frac{\lambda-(1+m)}{\lambda(1+m)}]$.

It remains to show the existence of α within this interval for which $A_\mu \succeq 0$. Since its first line and column were set to zero, we have to show the positive semidefiniteness only for the submatrix $\bar{A}_\mu \coloneqq (A_\mu)_{(2:3,2:3)}$, which reads

$$\bar{A}_\mu = \begin{bmatrix} A_{(2,2)} + 2sA_{(2,3)} + s^2 A_{(3,3)} & A_{(2,3)} + sA_{(3,3)} \\ \cdot & A_{(3,3)} \end{bmatrix}. \tag{5.22}$$

Note that the element $A_{(3,3)}$ and $\det \bar{A}_\mu$ are unchanged by the linear transformation T_μ. In Lemma 5.A.1 we show that $A_{(3,3)} \geq 0$ under the shrunk interval $\alpha \in [0, \bar{\alpha}]$, where $\bar{\alpha}$ has a well-defined expression. Then, in Lemma 5.A.2 we demonstrate that either $\alpha = 0$ or $\alpha = \alpha^*$, where α^* is the unconstrained maximizer of $\det \bar{A}_\mu(\alpha)$, are feasible solutions for which $\det \bar{A}_\mu(\alpha) \geq 0$. This concludes the proof for regime p_μ^∇.

Case 2. Regime p_L^∇: The PEP primal solution informs that

$$t_L \coloneqq \Delta g_u - \ell \Delta u = 0;$$

$$r_L \coloneqq \Delta g_{u,v} + \frac{m}{\ell}(1 - \frac{1+\ell}{\lambda})\Delta g_u = 0.$$

Similarly to the proof of regime p_μ^∇, we define the transformation

$$T_L \coloneqq \begin{bmatrix} 1 & 0 & 0 \\ \ell & 1 & 0 \\ sm & s\frac{m}{\ell} & 1 \end{bmatrix}, \tag{5.23}$$

where $s = \frac{1+\ell-\lambda}{\lambda}$, such that $x = T_L x_\mu^\nabla$, with $x_L^\nabla = \left[\Delta u,\, t_L,\, r_L\right]^\top$. Then inequality (5.21) is equivalent to $\tilde{x}^\top (T_L^\top A T_L)\tilde{x} \geq 0$ and zeroing $A_L^{(1,1)}$, $A_L^{(1,2)}$ and $A_L^{(1,3)}$ yields the system of equations:

$$\begin{cases} 0 = A_{(1,1)} + 2\ell A_{(1,2)} + 2smA_{(1,3)} + \ell^2 A_{(2,2)} + 2s\ell m A_{(2,3)} + s^2 m^2 A_{(3,3)}; \\ 0 = A_{(1,2)} + \dfrac{sm}{\ell} A_{(1,3)} + \ell A_{(2,2)} + 2smA_{(2,3)} + \dfrac{(sm)^2}{\ell} A_{(3,3)}; \\ 0 = A_{(1,3)} + \ell A_{(2,3)} + smA_{(3,3)}. \end{cases}$$

Substituting the expressions for the entries of matrix A and performing the simplifications, the system of equations becomes:

$$\begin{cases} 0 = \lambda(1-\ell)[2(1+\ell) - \lambda] - p[(1-\lambda)(\ell - m) + 1 - \ell m]^2; \\ 0 = \ell\lambda(2\ell - \lambda) + p[(1-\lambda)(\ell - m) - \ell m][(1-\lambda)(\ell - m) + 1 - \ell m] \\ \qquad - \alpha_{u_1,u_0}\ell\lambda^2 + \alpha_{v_0,u_1} m(1-\lambda)(1+\ell-\lambda); \\ 0 = \alpha_{v_0,u_1}(1+\ell-\lambda) - p[(1+\ell)(1-m) - \lambda(\ell - m)]. \end{cases}$$

The first equation yields $p = p_L^{\nabla}$, while from the other two we get:

$$\alpha_{u_1,u_0} = \frac{2(1-\lambda)}{\lambda} + \frac{(2-\lambda)[2\ell-(1+m)]}{D};$$
$$\alpha_{v_0,u_1} = \frac{D}{\bar{s}} p_L^{\nabla}, \tag{5.24}$$

where $\bar{s} \coloneqq 1+\ell-\lambda$ and $D \coloneqq 1-\ell m+(1-\lambda)(\ell-m)$.

The condition $\alpha_{u_1,u_0} \geq 0$ restricts the domain of regime p_L^{∇}. We have that α_{u_1,u_0} monotonically decreases with m since it holds that

$$\frac{d\alpha_{u_1,u_0}}{dm} = \frac{-(2-\lambda)}{\lambda} p_L^{\nabla} < 0.$$

Therefore, $\alpha_{u_1,u_0} \geq 0$ only if $m \leq \bar{m}$, where $\bar{m}$ is the unique root of α_{u_1,u_0}.

$$\bar{m} \coloneqq \frac{(2-\lambda)^2-2(1-\ell)}{(1-\lambda)^2+2\ell(1-\lambda)+1} \geq m,$$

This bound on m appears in the branch of p_L^{∇} from (5.11). Moreover, as shown in Lemma 5.A.6, this condition additionally implies $\bar{s} > 0$. Furthermore, $D = \bar{s}(1-m)+\lambda(1-\ell) > 0$ and hence we also have that $\alpha_{v_0,u_1} \geq 0$.

It remains to determine the existence of $\alpha \coloneqq \alpha_{u_0,u_1} \geq 0$ such that $\alpha_{u_0,v_0} \geq 0$, $\alpha_{u_1,v_0} \geq 0$ and $A_L \succeq 0$. Recall that

$$\alpha_{u_0,v_0} = 1+\alpha_{u_1,u_0}-\alpha_{u_0,u_1} \geq 0;$$
$$\alpha_{u_1,v_0} = \alpha_{u_0,u_1}-\alpha_{u_1,u_0}+\alpha_{v_0,u_1}-1 \geq 0,$$

implying $\alpha \in \left[[1+\alpha_{u_1,u_0}-\alpha_{v_0,u_1}\right]_+, 1+\alpha_{u_1,u_0}]$. This interval is non-empty due to nonnegativity of α_{u_1,u_0} and α_{v_0,u_1}.

The first line and column of A_L being set to zero, it is left to prove the existence of α such that the submatrix $\bar{A}_L \coloneqq (A_L)_{(2:3,2:3)}$ is positive semidefinite, where

$$\bar{A}_L = \begin{bmatrix} A_{(2,2)} + \frac{2ms}{\ell} A_{(2,3)} + (\frac{ms}{\ell})^2 A_{(3,3)} & A_{(2,3)} + \frac{ms}{\ell} A_{(3,3)} \\ \cdot & A_{(3,3)} \end{bmatrix}. \tag{5.25}$$

Similarly to the previous case, the element $A_{(3,3)}$ is unchanged when applying the linear transformation T_L and, after substitution, it reads as

$$A_{(3,3)} = p_L^{\nabla}[\ell-m+\frac{2(1-\ell^2)}{\bar{s}}],$$

whose positivity is dictated by $\bar{s} > 0$ (proved in Lemma 5.A.6).

In Lemma 5.A.7 we show there exists $\alpha \in [[1+\alpha_{u_1,u_0} - \alpha_{v_0,u_1}]_+ \, , \, 1+\alpha_{u_1,u_0}]$ such that $\det \bar{A}_L(\alpha) \geq 0$. As in the previous case, this determinant is exactly the second principal minor of the initial matrix A, namely $\det \bar{A}_L = A_{(2,2)}A_{(3,3)} - (A_{(2,3)})^2$.

Case 3. Regime p_3^∇. This regime is valid when both proofs for p_μ^∇ and p_L^∇ break. The PEP dual solution reveals the structure of non-zero multipliers, i.e., $\alpha_{u_0,v_0} = 1$, $\alpha_{u_1,v_0} = 1 + \alpha$ and $\alpha_{v_0,u_1} = 2 + \alpha$, with some $\alpha \geq 0$. The extreme (worst) case is characterized by $\det A(\alpha, p) = 0$ and solving it for p yields $p = p(\alpha)$.[2] The critical points of $p(\alpha)$ are $\frac{-2m}{1+m} \geq 0$ and $\frac{-2}{1+\ell} - 1 < 0$, therefore we select $\alpha = \frac{-2m}{1+m}$. Plugging it back in the expression of p gives

$$p_3^\nabla := p\Big(\frac{-2m}{1+m}\Big) = \frac{\lambda[2(1+\ell) - \lambda] - \frac{4(\ell-m)}{1-m}}{(1-\ell)(1+m)^2 - (\ell-m)(2-\lambda)^2}.$$

On the feasible domain of p_3^∇ we have $p_3^\nabla \in [0,1]$ (see Lemma 5.A.8). Finally, in Lemma 5.A.9 we prove that $A \succeq 0$ with the choice $\alpha = \frac{-2m}{1+m}$. $\square$

The current proof uses five multipliers, but the primal-PEP also shows that the sixth possible interpolation inequality for φ_1 holds exactly (due to the extremal worst-case construction). Therefore, there exists additional flexibility for proving the decrease.

The proofs of regimes p_μ^∇ and p_L^∇ are parametrized in $\alpha = \alpha_{u_0,u_1}$ which must stay in a definite range. For regime p_L^Δ, as shown in Lemma 5.A.7, there exists a unique multiplier that works on the entire feasible range, namely the one maximizing the minor of interest. However, for p_μ^∇, the maximizer of the minor is not always within the feasible range.

Moreover, the multiplier values coincide with the dual-PEP numerics only on the regime boundaries, where they are uniquely determined. Inside the domains of p_μ^∇ and p_L^∇, the SDP solver may return any feasible dual point, often without a stable pattern, so reverse-engineering multipliers from numerics is unreliable. This explains why PEP certificates can be hard to interpret and why one should avoid trying to identify all multipliers. Instead, provide an independent analytical proof and then check it for consistency with the PEP output. Nevertheless, PEP does confirm the mandatory relations among multipliers – those enforcing the linear combination of function values in the target inequality and those implied by cancellations in the slack matrix.

[2] For brevity, we omit to provide these expressions explicitly as they are very long; the interested reader can check the symbolic script accompanying the proof.

5.6 Numerical experiments

We test the optimized tunings from Proposition 5.3.1 and Conjecture 5.3.1 (for the *iterate progress*), and the tuning for the *best subgradient* obtained by numerically maximizing the denominator in Theorem 5.4.1. We consider a simple family: minimization of a quadratic function subject to unit ℓ_∞- or ℓ_2-ball constraints. Specifically,

$$\varphi_1(x) := \frac{1}{2}x^\top \Sigma x,$$

where Σ's eigenvalues are prescribed in the range $[\mu, L]$. Its proximal mapping is

$$\mathrm{prox}_{\gamma\varphi_1}(x) = (I + \gamma\Sigma)^{-1}\, x.$$

For the constraint term, we take either the indicator of the unit box (i.e., the ℓ_∞-ball) $\varphi_2 = \delta_{\mathbb{B}(0,1)}$ or of the unit Euclidean ball $\varphi_2 = \delta_{\mathbb{B}_2(0,1)}$. The corresponding proximal mappings are projections:

$$\mathrm{prox}_{\delta_{\bar{B}(0,1)}}(x) = x \odot \frac{1}{\max(|x|, \vec{1})} \quad \text{(projection onto unit box; elementwise)}$$

$$\mathrm{prox}_{\delta_{\bar{B}(0,1)}}(x) = \frac{x}{\max(\|x\|_2, 1)} \quad \text{(projection onto unit ball).}$$

Setup: We fix $L = 1$ and, for several values of μ, sample a grid over $\lambda \in (0, 2)$ and $\gamma \in (0, \min\{\frac{1}{L}, \frac{2-\lambda}{-2\mu}\}$. Each grid uses 52 samples per parameter and includes the optimized values of $(\gamma^*, \lambda^*)^\Delta$ and $(\gamma^*, \lambda^*)^\nabla$. We report the mean number of iterations to reach 10^{-8} accuracy in both criteria (iterate progress and residual subgradient), averaged over n random instances in dimension d. A run is terminated at $N_{\max} = 1000$ iterations; runs that do not meet the tolerance by $N_{\max}$ are excluded from the averages.

As a general remark, empirical convergence is often faster than the sublinear worst-case rates; nevertheless, there is a clear link between the optimized worst-case tunings and empirically good choices in practice.

Box constraints. In Table 5.6.1 we report the best-performing grid configuration alongside the performance of the optimized tunings for $n = 50$ instances in dimension $d = 20$. The optimized tunings typically yield larger gains as nonconvexity increases (more negative μ). Between the two worst-case measures, it is not clear which one is uniformly superior: optimizing the iterate progress gives slightly better performance as μ decreases, whereas minimizing the guaranteed residual subgradient performs better for μ closer to 0. Larger-scale experiments appear necessary to assess how consistently close they are to the empirical best tunings. Empirically, when $\mu > -L$, the optimal tuning is close to the largest possible stepsize, for which we take the maximum possible relaxation, namely $\gamma L = 1$ and $\lambda = 2(1 + \mu)$.

Figure 5.6.1 compares the empirical and theoretical performance in the case $\mu = -L = -1$. In Figure 5.6.1a, we report results for the case $\mu = -1$ on $n = 30$ random generated instances in dimension $d = 30$. The theoretically motivated tunings lie near the empirically optimal region but remain slower: the tuning based on the iterate progress bound requires roughly twice as many iterations, while the tuning based on the residual subgradient bound is almost three times slower. Nonetheless, both theoretical tunings fall within a regime of relatively fast convergence, each requiring fewer than 200 iterations on average.

Ball constraints. In Figure 5.6.2a, we report results for the case $\mu = -1$ on $n = 30$ random generated instances in dimension $d = 50$. To increase the difficulty of the problem, we use a *softmax* distribution of eigenvalues, so that they are clustered near the extreme curvature bounds μ and L. Figure 5.6.2 compares the empirical and theoretical performance in the case $\mu = -L = -1$.

The theoretically motivated tunings lie near the empirically optimal region but remain slower: the tuning based on the iterate progress bound requires 22% more iterations, while the tuning based on the residual subgradient bound is about 50% slower. Nonetheless, both theoretical tunings fall within a regime of relatively fast convergence, each requiring fewer than 50 iterations on average.

5.7 Non-invariance of DRS to curvature shifting

We conclude with a short note showing that *Douglas–Rachford splitting (DRS) is not invariant* to *curvature shifting*, in contrast with the behaviour we exploited for DCA. Consider

$$\min_{s \in \mathbb{R}^d} \varphi(s) = \varphi_1(s) + \varphi_2(s),$$

and fix a curvature shift $c \in \mathbb{R}$ and a stationary point $s^* \in \mathbb{R}^d$. Define the shifted pair

$$\tilde{\varphi}_1(s) = \varphi_1(s) - \frac{c}{2}\|s - s^*\|^2,$$

$$\tilde{\varphi}_2(s) = \varphi_2(s) + \frac{c}{2}\|s - s^*\|^2,$$

so that $\varphi_1(s) + \varphi_2(s) \equiv \tilde{\varphi}_1(s) + \tilde{\varphi}_2(s)$. We compare one DRS step from an initial point s in the original and shifted problems.

One DRS step. For (φ_1, φ_2) we have that:

$$\begin{cases} u = \operatorname{prox}_{\gamma\varphi_1}(s) \\ v = \operatorname{prox}_{\gamma\varphi_2}(2u - s) \\ s^+ = s + \lambda(v - u), \end{cases}$$

whereas for $(\tilde{\varphi}_1, \tilde{\varphi}_2)$ it holds that:

$$\begin{cases} \tilde{u} = \left[\mathrm{id} + \gamma\nabla\tilde{\varphi}_1\right]^{-1}(s) \\ \tilde{v} = \left[\mathrm{id} + \gamma\partial\tilde{\varphi}_2\right]^{-1}(2\tilde{u} - s) \\ \tilde{s}^+ = s + \lambda(\tilde{v} - \tilde{u}). \end{cases}$$

We ask whether $\tilde{s}^+ = s^+$ holds in general. The answer is negative.

Example 5.7.1 (Counterexample to shifting invariance for DRS)**.** *Let* $\varphi_1(s) = \frac{1}{2}s^\top As$ *and* $\varphi_2(s) = \frac{1}{2}s^\top As$ *with* $A = I$*, and take* $s^\star = 0$*. Then*

$$\operatorname{prox}_{\gamma\varphi_1}(x) = \operatorname{prox}_{\gamma\varphi_2}(x) = \frac{1}{1+\gamma}\,x.$$

Define the shifted pair with $c = -1$*:*

$$\tilde{\varphi}_1(s) = \tfrac{1}{2}s^\top As - \tfrac{c}{2}\|s\|^2 = \tfrac{1-c}{2}\|s\|^2, \qquad \tilde{\varphi}_2(s) = \tfrac{1}{2}s^\top As + \tfrac{c}{2}\|s\|^2 = \tfrac{1+c}{2}\|s\|^2,$$

so that

$$\operatorname{prox}_{\gamma\tilde{\varphi}_1}(x) = \frac{1}{1+\gamma(1-c)}\,x, \qquad \operatorname{prox}_{\gamma\tilde{\varphi}_2}(x) = \frac{1}{1+\gamma(1+c)}\,x.$$

With $\gamma = 0.5$*,* $\lambda = 1$*, and* $s = 0.5$ *we obtain:*

- *Original DRS:* $u = \frac{2}{3}s$*,* $v = \frac{2}{3}(2u - s) = \frac{2}{9}s$*, hence* $s^+ = s + (v - u) = \frac{5}{9}s = 0.277$*.*
- *Shifted DRS (*$c = -1$*):* $\operatorname{prox}_{\gamma\tilde{\varphi}_1}(x) = \frac{1}{1+2\gamma}x = \frac{1}{2}x$ *and* $\operatorname{prox}_{\gamma\tilde{\varphi}_2}(x) = \frac{1}{1+0}x = \mathrm{id}$*, so* $\tilde{u} = \frac{1}{2}s$*,* $\tilde{v} = 2\tilde{u} - s = 0$*, and* $\tilde{s}^+ = s - (\frac{1}{2}s) = \frac{1}{2}s = 0.25$*.*

Thus $\tilde{s}^+ \neq s^+$ *(here,* $0.25 \neq 0.277$*), although the objective* $\tilde{\varphi}_1 + \tilde{\varphi}_2$ *equals* $\varphi_1 + \varphi_2$*.*

Unlike PGD or DCA, the DRS iteration depends on how curvature is split between the two terms: **DRS is not invariant** to curvature shifting (unless in degenerated cases such as $c = 0$, or when s is already a fixed point).

With respect to the worst-case bounds for the iterate progress and the residual subgradient, one can optimize the curvature shifting parameter. The tight stepsize bounds that ensure convergence become

$$\gamma \in \left(0, \min\left\{\frac{1}{L-c}, \frac{2-\lambda}{-2(\mu-c)}\right\}\right),$$

which requires c to lie in the interval

$$c \in \left(\mu + \frac{2-\lambda}{2\gamma}, L - \frac{1}{\gamma}\right).$$

Jointly optimizing the three parameters: stepsize γ, relaxation λ, and curvature shifting c, is left for future work. A key point is to define a suitable way to compare the different progress guarantees across different choices, since the resulting DRE value depends on both the stepsize and the curvature parameters.

5.8 Conclusion

We developed tight convergence guarantees for the Douglas–Rachford Splitting (DRS) method when the objective decomposes into a smooth term and a lower semicontinuous term with an existing proximal operator for sufficiently small stepsizes. Using the Douglas–Rachford Envelope (DRE) as a Lyapunov function, we established bounds on the *iterate progress* and the *residual subgradient*, which we conjecture are exact.

Our proofs are fully analytical yet informed by PEP insights, yielding a transparent recipe for constructing the bounds and selecting the associated multipliers. Notably, a third, non–extreme curvature regime emerges for the iterate progress metric when $\lambda \in (1, 2)$, while an analogous regime appears for the subgradient metric when $\lambda \in (0, 1)$. This raises a natural question: are these standard metrics optimal for worst-case analysis, or do they introduce artifacts tied to worst-case constructions?

Practically, our results enable principled tuning of the DRS stepsize and relaxation parameters to optimize worst-case behaviour for both metrics. Finally, didactic numerical examples illustrate that these tunings translate into good empirical performance.

Table 5.6.1: Experiment: Smooth nonconvex quadratic s.t. unit box constraints. We report the mean iteration count to reach 10^{-8} accuracy on problems with $d = 20$, averaged over 50 random initializations. The parameters γ and λ are varied over their admissible ranges on a 52×52 grid. We test several values of the lower curvature bound μ (with $L = 1$). For each criterion (iterate progress and best subgradient) we report the percentile rank (top %) of the theory-recommended tuning within the grid and also display the empirically best configurations. Runs exceeding 1000 iterations are terminated and excluded from the aggregate statistics.

μ	**Optimized measure**	γ	λ	N^*	**Percentile**
0	Iterate progress (Δ)	0.5000	1.5000	91.9	33.73
	Residual subgradient (∇)	0.8602	1.3379	59.2	14.78
	Experiment	0.9990	1.9633	31.0	–
-0.1	Iterate progress (Δ)	0.5000	1.5000	432.7	39.26
	Residual subgradient (∇)	0.8068	1.2511	343.9	25.65
	Experiment	0.9403	1.6422	137.7	–
$-\frac{2}{7}$	Iterate progress (Δ)	0.5000	1.5000	117.4	22.93
	Residual subgradient (∇)	0.6956	1.1089	115.4	21.87
	Experiment	0.9794	1.1089	42.5	–
$-\frac{1}{3}$	Iterate progress (Δ)	0.4963	1.4635	144.4	25.73
	Residual subgradient (∇)	0.6672	1.0783	152.5	28.7
	Experiment	0.9599	1.3211	53.9	–
-0.5	Iterate progress (Δ)	0.4797	1.3411	87.2	19.98
	Residual subgradient (∇)	0.5758	0.9900	85.7	19.32
	Experiment	0.9786	0.9630	43.1	–
$-\frac{2}{3}$	Iterate progress (Δ)	0.4582	1.2303	148.3	18.33
	Residual subgradient (∇)	0.5000	0.9260	207.4	35.13
	Experiment	0.6528	1.0740	68.8	–
-1	Iterate progress (Δ)	0.4064	1.0550	71.1	6.12
	Residual subgradient (∇)	0.3899	0.8436	118.3	24.47
	Experiment	0.4706	0.8930	40.5	–
-2	Iterate progress (Δ)	0.2704	0.8114	40.6	4.89
	Residual subgradient (∇)	0.2300	0.7389	52.1	15.41
	Experiment	0.2744	0.6432	29.8	–

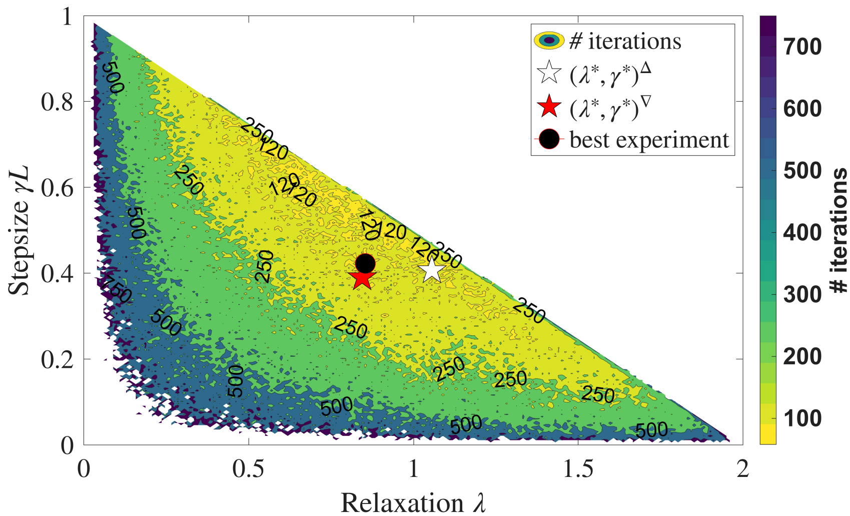

(a) Contour plot of the average number of iterations to reach 10^{-8} accuracy (mean over 30 runs with randomly generated initial points) for a problem of dimension $d = 50$ in the smooth nonconvex setting ($\mu = -L$) with *linearly distributed eigenvalues*, subject to *unit box constraints*. Highlighted markers indicate the following parameter choices: (i) best empirical tuning: $(\gamma^*, \lambda^*, N^*) = (0.4223, 0.8531, 58)$; (ii) theoretical tuning from the subgradient bound: $(\gamma^*, \lambda^*, N^*)^\nabla = (0.3899, 0.8436, 154.7)$, ranked within the top 16.31% of all tested parameter pairs converging in less than 1000 iterations; (iii) theoretical tuning from the iterate progress bound: $(\gamma^*, \lambda^*, N^*)^\Delta = (0.4064, 1.0550, 117.7)$, ranked within the top 6.15%.

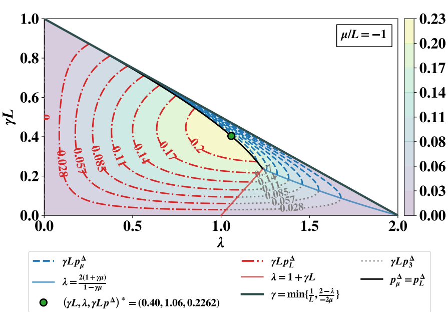

(b) Worst-case iterate progress

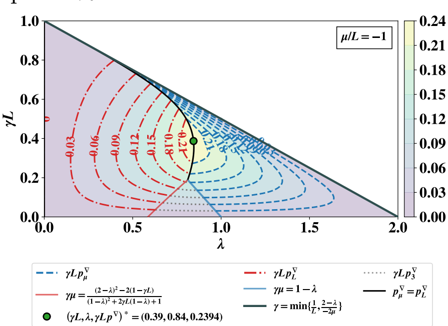

(c) Worst-case residual subgradient

Figure 5.6.1: (DRS) Comparison of empirical results with the theoretical worst-case for smooth functions with $\mu = -L$, under both performance criteria, using *linearly* distributed eigenvalues and a **box-constrained** setup.

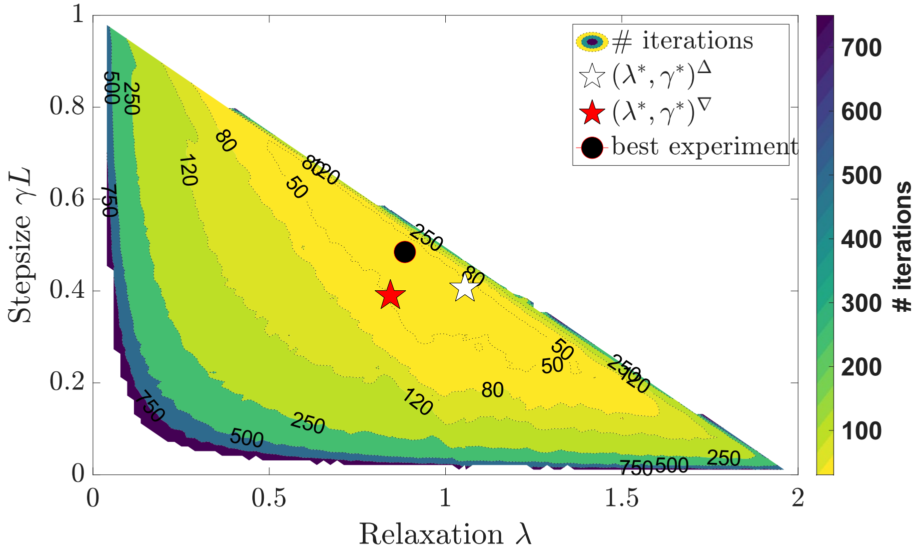

(a) Contour plot of the average number of iterations to reach 10^{-8} accuracy (mean over 30 runs with randomly generated initial points) for a problem of dimension $d = 50$ in the smooth nonconvex setting ($\mu = -L$) with *softmax-like distributed eigenvalues*, subject to *unit ball constraints*. Highlighted markers indicate the following parameter choices: (i) best empirical tuning: $(\gamma^*, \lambda^*, N^*) = (0.4849, 0.8847, 30.6)$; (ii) theoretical tuning from the subgradient bound: $(\gamma^*, \lambda^*, N^*)^{\nabla} = (0.3899, 0.8436, 47.6)$, ranked within the top 10.57% of all tested parameter pairs converging in less than 1000 iterations; (iii) theoretical tuning from the iterate progress bound: $(\gamma^*, \lambda^*, N^*)^{\Delta} = (0.4064, 1.0550, 37.5)$, ranked within the top 2.75%.

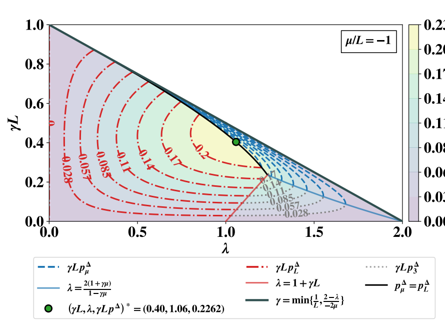

(b) Worst-case iterate progress

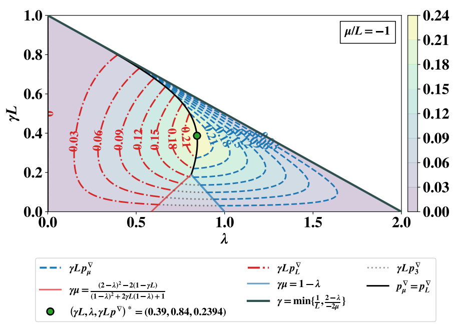

(c) Worst-case residual subgradient

Figure 5.6.2: Comparison of empirical results with the theoretical worst-case for smooth functions with $\mu = -L$, for both performance criteria, using a *softmax-like* distributed eigenvalue spectrum and a **ball-constrained** setup.

Appendix

5.A Helper lemmas for deriving the bounds on residual subgradient

In this appendix, we provide the lemmas used in the proof of Theorem 5.4.1, given in Section 5.5.5, showing the descent bound on residual subgradient.

5.A.1 Helper lemmas for regime p_μ^∇

In this section, we consider the feasible range of regime p_μ^∇, with $\ell \in (0,1)$, $m \leq 0$, $\lambda \in [1+m, 2(1+m)) \cap (0,2)$ and $S = p_\mu^\nabla(\ell,m,\lambda) - p_L^\nabla(\ell,m,\lambda) \leq 0$, where p_μ^∇ and p_L^∇ are defined in (p_μ^∇) and (p_L^∇), respectively. A handy observation is that $1+m > 0$, since $\lambda > 0$.

Lemma 5.A.1 (Positivity of $A_{(3,3)}$)**.** *On the feasible range of regime p_μ^∇ it holds that $A_{(3,3)}(\alpha) \geq 0$, with $\alpha \geq 0$, if and only if $\alpha \leq \bar{\alpha}$, where*

$$\bar{\alpha} \coloneqq \frac{[\lambda - (1+m)][2(1+m) - \lambda]\left[\frac{4(1-m^2)}{\lambda[2(1+m)-\lambda]} - (\ell - m)\right]}{2(1-m)(1+m)^2}. \tag{5.26}$$

Moreover, $\bar{\alpha} \in [0, 2\frac{\lambda-(1+m)}{\lambda(1+m)}]$.

Proof. Substituting into $A_{(3,3)} = 2\alpha_{v_0,u_1} - (\ell - m)p$, after simplifications we get

$$\frac{A_{(3,3)}}{p_\mu^\nabla} = \frac{-2(1-m)(1+m)^2}{[\lambda - (1+m)][2(1+m) - \lambda]}\alpha +$$

$$\frac{(1-m)[2\lambda + (2-\lambda)[2(1+m) - \lambda]]}{\lambda[2(1+m) - \lambda]} + 1 - \ell.$$

Imposing its nonnegativity implies the constraint $\alpha \leq \bar{\alpha}$. This bound is nonnegative, as implied by lower bounding with $\ell \leq 1$, and is actually tighter than $2\frac{\lambda-(1+m)}{\lambda(1+m)}$ due to

$$2\frac{\lambda-(1+m)}{\lambda(1+m)} = \bar{\alpha} + \frac{(\ell-m)[\lambda-(1+m)][2(1+m)-\lambda]}{2(1-m)(1+m)^2}.$$

□

Lemma 5.A.2 (Positivity of $\det \bar{A}_\mu$)**.** *On the feasible range of regime* p_μ^∇, *there exists* $\alpha \in [0, \bar{\alpha}]$, *with* $\bar{\alpha}$ *given in* (5.26), *for which* $\det \bar{A}_\mu(\alpha) \geq 0$, *where* $\bar{A}_\mu(\alpha)$ *is defined in* (5.22).

Proof. We expand the expression of $\det \bar{A}_\mu(\alpha)$ from (5.22) and observe that it is a concave quadratic in parameter α:

$$\det \bar{A}_\mu(\alpha) = \frac{a_2\alpha^2 + a_1\alpha + a_0}{b},$$

with

$$\begin{aligned} b &= \lambda^2(1-m^2)(1+m)^2[\lambda-(1+m)]^2 \geq 0; \\ a_2 &= -\lambda^2(1-m)(1+m)^3[\lambda-(1+\ell)]^2 \leq 0; \\ a_1 &= 2\lambda(1+\ell)(1+m)[\lambda-(1+m)]q_1; \\ a_0 &= (1+m-\lambda)^2 q_0, \end{aligned} \tag{5.27}$$

where q_0 and q_1 are polynomials in (ℓ, m, λ):

$$\begin{aligned} q_0(\ell,m,\lambda) = &-2(1-m)[2+2m(1-\lambda)^2-\lambda(2-\lambda)^2]\,\ell^2 \\ &-(1+m)[8-\lambda(2+\lambda)(2-\lambda)^2+2m(\lambda^3-2\lambda^2+4\lambda-4)]\,\ell + \\ &m^2(2\lambda^3-4\lambda^2+4)-m\lambda^2(\lambda^2-6\lambda+4)-\lambda^4+8\lambda^2-4; \end{aligned} \tag{5.28}$$

$$\begin{aligned} q_1(\ell,m,\lambda) = &[2(1-m^2)(1-\lambda)+\lambda^2(3+m-\lambda)]\ell + \\ &\lambda^3-[2+m(1+m)]\lambda^2+2(1-m^2)(1-\lambda). \end{aligned} \tag{5.29}$$

We examine the behaviour of $\det \bar{A}_\mu(\alpha)$ at the extremal points of the interval and the unconstrained optimizer, namely $\alpha \in \{0, \bar{\alpha}, \alpha^*\}$, where $\alpha^* := \arg\max_\alpha[\det \bar{A}_\mu(\alpha)]$.

At the lower bound $\alpha = 0$ we obtain:

$$\det \bar{A}_\mu(\alpha = 0) = \frac{(1+m-\lambda)^2}{b} q_0,$$

hence $\operatorname{sgn}(\det \bar{A}_\mu(\alpha = 0)) = \operatorname{sgn}(q_0)$. On the other hand, at the upper bound the determinant $\det \bar{A}_\mu(\bar{\alpha})$ is negative since it can be written as:

$$\det \bar{A}_\mu(\bar{\alpha}) = \frac{-\lambda^2[(1-\lambda)(\ell-m) - \ell(\ell+m) + 2]^2(1+m-\lambda)^2[2(1+m)-\lambda]^2}{4(1-m)(1+m)b}.$$

Moreover, the slope of $\det \bar{A}_\mu(\alpha)$ is negative at $\bar{\alpha}$ (see Lemma 5.A.3):

$$(\det \bar{A}_\mu)'(\alpha = \bar{a}) = \frac{2a_2\bar{a} + a_1}{b} \leq 0,$$

hence $\bar{\alpha} \geq \alpha^*$. Nonetheless, the unconstrained maximizer is

$$\alpha^* = \frac{a_1}{-2a_2} = \frac{(1+\ell)[\lambda - (1+m)]}{\lambda(1-m)(1+m)^2(1+\ell-\lambda)^2} q_1.$$

Evaluating its determinant we get

$$\det \bar{A}_\mu(\alpha^*) = \frac{a_1^2 - 4a_0a_2}{-4a_2b} = \frac{K(-S)}{-4a_2b},$$

where $K(\ell, m, \lambda)$ is defined as

$$K = \frac{4\lambda^3(1-m)(1+m)^4(1-\ell)[\lambda-(1+m)]^2[(1-\lambda)(\ell-m)+1-\ell m]^2}{\ell - m}\bar{K},$$

$$\bar{K} = 2(\ell+m)(1+\ell)(1+m) - \lambda(\ell^2 + \ell m + m^2 + \ell + m - 1).$$

We have that $\operatorname{sgn}(\alpha^*) = \operatorname{sgn}(q_1)$ and $\det \bar{A}_\mu(\alpha^*) \geq 0$ on the domain with $S \leq 0$, due to $\bar{K} \geq 0$ (see Lemma 5.A.4).

The existence of a feasible $\alpha \geq 0$ results from Lemma 5.A.5, stating that q_1 and q_0 cannot be simultaneously negative on the domain of regime p_μ^∇. Then, since we already have $\alpha^* \leq \bar{\alpha}$, it results that there exists is a feasible point $\alpha \in [0, \bar{\alpha}]$ with $\det \bar{A}_\mu(\alpha) \geq 0$, concluding the proof. □

Lemma 5.A.3. *On the feasible range of regime* p_μ^∇ *it holds that* $2a_2\bar{a} + a_1 \leq 0$, *where* $\bar{a}$ *is defined in* (5.26), *and* a_1 *and* a_2 *are given in* (5.27).

Proof. We have that

$$2a_2\bar{a} + a_1 = \lambda^2(1+m)[\lambda - (1+m)]G,$$

with

$$\begin{aligned} G \coloneqq\ & \ell^2(\ell+m)\left[2(1+m)-\lambda\right] \\ & -\ell\left[2(2-\lambda)m^2+2(1-\lambda)m+\lambda^3-4\lambda^2+3\lambda-2\right] \\ & +\left[-2(\lambda^2+3\lambda-3)m^2+(\lambda^3-4\lambda^2+3\lambda-2)m+2(\lambda^2-4\lambda+2)\right]. \end{aligned}$$

Recall $S = p_\mu^\nabla - p_L^\nabla$, which rewrites as

$$S = \frac{\lambda(\ell-m)}{[\ell m+(\ell-m)(\lambda-1)-1]^2(1-m)(1+m)^2}\bar{S},$$

where its sign is dictated by $\bar{S}$ defined as follows

$$\begin{aligned} \bar{S} \coloneqq & (4+4\ell-7\lambda+2\lambda^2)+\ell\lambda(4\lambda-\lambda^2-5)+ \\ & +m(\lambda^3-4\lambda^2+2\lambda)+m^2(3\ell\lambda-4\ell+5\lambda-2\lambda^2-4). \end{aligned}$$

We show that $G \le 0$ whenever $\bar{S} \le 0$, by proving that $H \coloneqq \frac{-(G-S)}{1-\ell}$ is nonnegative. After simplifications, H has an affine dependence on λ:

$$H = [1-\ell^2-(\ell+m)(1+m)]\lambda + 2(1+\ell)(1+m)(\ell+m).$$

Its minimum is reached at one of the edges of the λ-interval $[1+m, 2(1+m)]$. At the higher end,

$$H\big(\lambda = 2(1+m)\big) = 2(1-m)(1+m)^2 \ge 0.$$

At the lower end,

$$H\big(\lambda = 1+m\big) = (1+m)Z,$$

with

$$Z(\ell,m) = (1+m)(1+\ell)+\ell^2-m^2.$$

We demonstrate that $Z \ge 0$ when $\bar{S} \le 0$. Note that it is convex quadratic in ℓ and also strictly increasing with ℓ since $\frac{dZ(\ell,m)}{d\ell} = 2\ell+(1+m) > 0$. Thus,

$$Z(\ell,m) \ge Z(\ell=0,m) = 1+m-m^2. \tag{5.30}$$

Because $1+m-m^2$ is increasing on $(-\infty,0]$, its minimum is reached for the minimum m satisfying $\bar{S}(m) \le 0$. This minimum, larger than -1 due to $1+m > 0$, is determined below.

With $\lambda = 1+m$ in the expression of $\bar{S}$ we get

$$\bar{S}(\lambda = 1+m) = (1+m)[2\ell(1+m^2)-\psi(m)],$$

with $\psi(m) = m^3 - m^2 + 3m + 1$. Due to $\bar{S} \le 0$, we have that $\psi(m) \ge 2\ell(1+m^2) > 0$. Moreover, $\psi'(m) = 3m^2 - 2m + 3 = 3(m - \frac{1}{3})^2 + \frac{8}{3} > 0$, hence $\psi(m)$ is strictly increasing. Together, we have a lower bound for $m \in (-1, 0]$, namely the scalar r for which $\psi(r) = 0$. Because no exact root is available in analytical form, we use an under approximation of it. More precisely, due to $\psi(\frac{-1}{3}) = \frac{-4}{27}$ and $\psi(0) = 1$, we have that $r \in (\frac{-1}{3}, 0)$ and we take $\bar{r} := \frac{-1}{3} \le r$. Finally, by replacing in (5.30) we obtain

$$Z(\ell, m) \ge Z(\ell = 0, m) = 1 + m - m^2 \ge 1 + \bar{r} - \bar{r}^2 = \frac{5}{9} > 0.$$

Consequently, the edge $H(\lambda = 1 + m)$ is also positive, hence $G < S$. □

Lemma 5.A.4 (Proof of $\bar{K} \ge 0$)**.** *On the feasible range of regime p_μ^∇ it holds that $\bar{K}(\gamma L, \gamma\mu, \lambda) \ge 0$, where*

$$\bar{K}(\ell, m, \lambda) := 2(\ell + m)(1 + \ell)(1 + m) - \lambda(\ell^2 + \ell m + m^2 + \ell + m - 1).$$

Proof. The first term is always positive. If $\ell^2 + \ell m + m^2 + \ell + m - 1 \le 0$, then the result is trivial. Otherwise, we lower bound $\bar{K}$ using that $\lambda < 2$:

$$\bar{K} > 2[m^2\ell + m\ell(1 + \ell) + 1].$$

This lower bound is a convex quadratic in $m \le 0$, with minimizer $m = \frac{-(1+\ell)}{2}$ and minimum $\frac{(1-\ell)(\ell^2+3\ell+4)}{4} > 0$. □

Lemma 5.A.5. *On the feasible range of regime p_μ^∇ it holds that*

$$\{(\ell, m, \lambda) : q_0(\ell, m, \lambda) < 0\} \cap \{(\ell, m, \lambda) : q_1(\ell, m, \lambda) < 0\} = \emptyset,$$

where q_0 and q_1 are defined in (5.28) and (5.29), respectively.

Proof. Blueprint of the proof:

i) Firstly, we show that q_0 is a concave quadratic function in variable ℓ and strictly decreasing for $\ell \in (0, 1)$. For this, we prove that $d^2 q_0(\ell, m, \lambda)/d\ell^2 < 0$ and $dq_0(\ell, m, \lambda)/d\ell|_{\ell=0} < 0$.

ii) Then we demonstrate that $q_1(\ell, m, \lambda) \le 0$ for all $\ell \in (0, \underline{\ell}]$, with a well-defined $\underline{\ell}$.

iii) Finally, we show that $q_0(\underline{\ell}, m, \lambda) \ge 0$, thus the conclusion.

Reinterpreting the interval of λ we have that $m \in [\frac{\lambda}{2} - 1, [\lambda - 1]_-]$.

i) We rewrite $q_0(\ell, m, \lambda) = -2(1-m)r_2(m,\lambda)\ell^2 - (1+m)r_1(m,\lambda)\ell + r_0(m,\lambda)$, with

$$r_2(m,\lambda) = 2 + 2m(1-\lambda)^2 - \lambda(2-\lambda)^2$$

$$r_1(m,\lambda) = 8 - \lambda(2+\lambda)(2-\lambda)^2 + 2m(\lambda^3 - 2\lambda^2 + 4\lambda - 4)$$

$$r_0(m,\lambda) = m^2(2\lambda^3 - 4\lambda^2 + 4) - m\lambda^2(\lambda^2 - 6\lambda + 4) - \lambda^4 + 8\lambda^2 - 4.$$

(a) Proof of $d^2 q_0(\ell, m, \lambda)/d\ell^2 < 0$. It is equivalent to showing that $r_2(m,\lambda) \geq 0$. Because it has positive slope in m, its minimum is reached at the lower end of the m-interval, $m = \frac{\lambda}{2} - 1$. After substitution and simplifications we get $r_2(\frac{\lambda}{2} - 1, \lambda) = \lambda > 0$.

(b) Proof of $dq_0(\ell, m, \lambda)/d\ell|_{\ell=0} < 0$. It is equivalent to showing that $r_1(m,\lambda) \geq 0$, which is affine in m, thus its minimum is reached at one of the edges of the m-interval. Evaluating at these points we have:

$$r_1(m = \frac{\lambda}{2} - 1, \lambda) = 2(4-\lambda)[1 + (1-\lambda)^2] > 0;$$

$$r_1(m = \lambda - 1, \lambda) = (1-\lambda)^4 + 10(1-\lambda)^2 + 5 > 0;$$

$$r_1(m = 0, \lambda) = 8 - \lambda(2+\lambda)(2-\lambda)^2 > 8 - 2(2+2)(2-1)^2 = 0.$$

Since $m \leq \min\{0, \lambda - 1\}$ implies $\lambda \in [1, 2)$, the last inequality holds when $m = 0$.

ii) Let $w_1(m,\lambda) := -2(1-m^2)(\lambda - 1) + \lambda^2(3 + m - \lambda)$ be the slope in parameter ℓ of $q_1(\ell, m, \lambda)$; we prove that it is positive. For $\lambda \in (0, 1]$, it is readily seen that $w_1(m,\lambda) \geq 0$ since $m \geq -1$. When $\lambda \in (1, 2)$, we investigate the minimum value of w_1, which is convex quadratic in m, with minimum attained either at the critical point $m_w := \frac{-\lambda^2}{4(\lambda-1)}$ or at the endpoints of the m-interval $[\frac{\lambda}{2} - 1, 0]$. The critical point is not feasible because

$$m_w - \big(\frac{\lambda}{2} - 1\big) = \frac{-[1 + 3(1-\lambda)^2]}{4(\lambda - 1)} < 0.$$

Evaluating w_1 at the edges of the m-interval yields:

$$w_1\big(\frac{\lambda}{2} - 1, \lambda\big) = 2 - \frac{(2-\lambda)^2}{2} > 0;$$

$$w_1(0, \lambda) = 1 - (\lambda - 1)^3 + \lambda > 0.$$

Hence $q_1(\ell, m, \lambda)$ is monotonically increasing with ℓ and $q_1(\ell, m, \lambda) \leq 0$ for

$$\ell \leq \underline{\ell} := 1 - \frac{(1-m^2)(2-\lambda)^2}{w_1(m,\lambda)} < 1.$$

Additionally, we directly observe that $q_1(\ell = 1, m, \lambda) > 0$, using the rewriting

$$q_1(\ell, m, \lambda) = (1 - m^2)(2 - \lambda)^2 - (1 - \ell)w_1.$$

iii) Evaluating $q_0(\ell = \underline{\ell}, u, \lambda)$, after simplifications we get

$$q_0(\ell = \underline{\ell}, u, \lambda) = \frac{\lambda^2(2-\lambda)^2(1-m^2)[\lambda - (1+m)]}{[\lambda^2(3+u-\lambda) + 2(1-m^2)(1-\lambda)]^2} H(m, \lambda),$$

and we show the positivity of H which reads as

$$H(m, \lambda) := 4(\lambda - 1)m^3 - 2(\lambda^3 - 5\lambda^2 + 4\lambda - 2)m^2 + \\ (\lambda^4 - 10\lambda^3 + 24\lambda^2 - 12\lambda + 4)m + (\lambda^4 - 4\lambda^3 - 2\lambda^2 + 16\lambda - 4).$$

Using that:

$$H_0 := H(m, \lambda = 1 + m) = (1 + m)^2[7 - m(m^2 - 5u + 3)] \geq 0;$$

$$H_1 := H(m, \lambda = 2(1 + m)) = 4(1 - m)^2(1 + m) \geq 0,$$

we can rewrite H as follows:

$$H(m, \lambda) = \frac{2(1+m) - \lambda}{(1+m)} H_0 + \frac{\lambda - (1+m)}{(1+m)} H_1 + \\ [2(1+m) - \lambda][\lambda - (1+m)]R(m, \lambda),$$

where

$$R(m, \lambda) = -(1 + m)\lambda^2 + (-m^2 + 4m + 1)\lambda - m^3 + 5m^2 - 3m + 7.$$

Because $\lambda \in [1 + m, 2(1 + m)]$, it remains to prove $R(m, \lambda) \geq 0$ on the same interval. Note that $R(m, \lambda)$ is concave in λ, thus its minimum is reached at the boundaries:

$$R\big(m, \lambda = 1 + m\big) = 7 + (-m)(3m^2 - 5m + 1) \geq 0;$$

$$R\big(m, \lambda = 2(1 + m)\big) = 5 + (-m)[7m^2 + (1 + m) + 4] \geq 0.$$

In conclusion, $H(m, \lambda) \geq 0$ and consequently $q_0(\ell = \underline{\ell}, m, \lambda) \geq 0$, for all ℓ such that $q_1(\ell, m, \lambda) < 0$.

$\square$

5.A.2 Helper lemmas for regime p_L^∇

We consider the feasible range of regime p_L^∇, with $\ell \in (0,1)$, $m \leq 0$, $\lambda \in (0,2)$, $\lambda < 2(1+m)$, $m \leq \bar{m} := \frac{(2-\lambda)^2-2(1-\ell)}{(1-\lambda)^2+2\ell(1-\lambda)+1}$ and $S = p_\mu^\nabla(\ell,m,\lambda) - p_L^\nabla(\ell,m,\lambda) \geq 0$, where p_μ^∇ and p_L^∇ are defined in (p_μ^∇) and (p_L^∇), respectively.

Lemma 5.A.6. *On the domain of p_L^∇, with $m \leq \bar{m} := \frac{(2-\lambda)^2-2(1-\ell)}{(1-\lambda)^2+2\ell(1-\lambda)+1}$, it holds that $\bar{s} := 1+\ell-\lambda > 0$.*

Proof. Note that $\bar{s} > 0$ when $\lambda \leq 1$. Further on, we consider $\lambda \in (1,2)$. We prove $\bar{s} > 0$ by contradiction, namely we assume $\ell \leq \lambda - 1$. We upper bound the inequality $\lambda < 2(1+m)$ by using the condition

$$m \leq \bar{m} := \frac{(2-\lambda)^2 - 2(1-\ell)}{(1-\lambda)^2 + 2\ell(1-\lambda) + 1},$$

whose denominator is always positive because: (i) it is linear in ℓ; (ii) for $\ell = 0$ it becomes $1 + (1-\lambda)^2 > 0$; (iii) for $\ell = 1$ it reads $(2-\lambda)^2 > 0$. Then we obtain $\lambda < 2(1+\bar{m})$, which after multiplying by the denominator of $\bar{m}$ implies

$$0 \leq 2\big[\big(\lambda - \frac{3}{2}\big)^2 + \frac{7}{4}\big]\ell - \lambda^3 + 6\lambda^2 - 14\lambda + 8.$$

The r.h.s. is affine and increasing in ℓ. Then its maximum is reached at $\ell = \lambda - 1$ and the inequality becomes

$$0 \leq 2\big[\big(\lambda - \frac{3}{2}\big)^2 + \frac{7}{4}\big](\lambda - 1) - \lambda^3 + 6\lambda^2 - 14\lambda + 8 = -(2-\lambda)\lambda^2 < 0,$$

which is a contradiction! □

Lemma 5.A.7. *On the domain of regime p_L^∇, with α_{u_1,u_0} and α_{v_0,u_1} given in (5.24), there exists $\alpha \in \big[[1+\alpha_{u_1,u_0} - \alpha_{v_0,u_1}]_+\,,\ 1+\alpha_{u_1,u_0}\big]$ such that $\det \bar{A}_L(\alpha) \geq 0$, where $\bar{A}_L(\alpha)$ is defined in (5.25). Moreover, $\alpha^* := \arg\max_\alpha \bar{A}_L(\alpha)$ is a viable candidate solution.*

Proof. Recall the previously defined expressions: $\bar{s} = 1+\ell-\lambda$ (proved to be nonnegative in Lemma 5.A.6) and $D = 1 - \ell m + (1-\lambda)(\ell - m) \geq 0$.

Substituting the expressions of $p = p_L^\nabla$, α_{u_1,u_0} and α_{v_0,u_1}, we have that $\det \bar{A}_L(\alpha) = A_{(2,2)}A_{(3,3)} - (A_{(2,3)})^2$ is a concave quadratic in parameter α:

$$\det \bar{A}_L(\alpha) = \frac{a_2\alpha^2 + a_1\alpha + a_0}{b},$$

with $b = \lambda^2 D^3 \bar{s}^2 > 0$, $a_2 = -b < 0$,

$$a_1 = 2D^3\lambda(2-\lambda)(1+\ell)\, q_1;$$

$$q_1 = \bar{s}[1-(1-m)p_L^{\nabla}].$$

a_0 has a long algebraic expression which we do not present here explicitly. Using symbolic toolboxes, we are able to compute $\Delta \coloneqq a_1^2 - 4a_0a_2$,

$$\Delta = Sp_L^{\nabla}\,\frac{4\lambda^2\bar{s}^3D^6(1-m^2)^2(D+1-\ell^2)}{\ell-m}.$$

(From this expression, a compact formula of a_0 can be derived.)
Observe that Δ has the sign of $S \geq 0$. The determinant is maximized by

$$\alpha^* \coloneqq \frac{a_1}{-a_2}\frac{(1+\ell)(2-\lambda)}{\lambda\bar{s}^2}q_1$$

and the maximum value is

$$\det A_L(\alpha^*) = \frac{\Delta}{-4a_2b} = Sp_L^{\nabla}\,\frac{(1-m^2)^2(D+1-\ell^2)}{(\ell-m)\lambda^2\bar{s}}.$$

Since $S \geq 0$, it follows that $\det A_L(\alpha^*) \geq 0$. Let $\underline{\alpha} \coloneqq 1+\alpha_{u_1,u_0}-\alpha_{v_0,u_1}$ and $\bar{\alpha} \coloneqq 1+\alpha_{u_1,u_0} \geq 0$ be the bounds of the admissible interval. The expressions of these points are:

$$\underline{\alpha} = \frac{-2p_L^{\nabla}D+(1+\ell)(2-\lambda)}{\lambda\bar{s}};$$

$$\bar{\alpha} = \frac{(2-\lambda)(1+\ell-p_L^{\nabla}D)}{\lambda\bar{s}},$$

implying

$$\alpha^*-\underline{\alpha} = \frac{\lambda(1-\ell)^2(1+m)[2(1+\ell)-\lambda]}{\bar{s}D^2} = \frac{(1-\ell)(1+m)}{\bar{s}}p_L^{\nabla} \geq 0;$$

$$\bar{\alpha}-\alpha^* = \frac{\lambda(\ell-m)(1-\ell)(2-\lambda)[2(1+\ell)-\lambda]}{\bar{s}D^2} = \frac{(2-\lambda)(\ell-m)}{\bar{s}}p_L^{\nabla} \geq 0.$$

Therefore, $\alpha^* \in [\underline{\alpha},\bar{\alpha}]$. Finally, we show that α^*, which has the sign of q_1, is nonnegative. Using the rewriting $p_\mu^{\nabla} = \frac{\lambda}{1-m^2}(2-\frac{\lambda}{1+m})$, q_1 rewrites as

$$q_1 = \bar{s}[1-(1-m)p_\mu^{\nabla}+(1-m)(p_\mu^{\nabla}-p_L^{\nabla})] = \bar{s}\big[\frac{(1+m-\lambda)^2}{(1+m)^2}+(1-m)S\big],$$

which is nonnegative on the feasible domain with $S \geq 0$. To conclude, α^* is a choice that works within the feasible range. □

5.A.3 Helper lemmas for regime p_3^∇

Consider the feasible range of regime p_3^∇, which is complementary to the ranges of p_μ^∇ and p_L^∇. Specifically, we have that $\ell \in (0,1)$, $m \leq 0$, $\lambda \in (0, 1+m)$ and $m \geq \bar{m} := \frac{(2-\lambda)^2 - 2(1-\ell)}{(1-\lambda)^2 + 2\ell(1-\lambda)+1}$. The condition $\lambda \leq 1+m$ imposes $\lambda \leq 1$. Moreover, the constraint $0 \geq m \geq \bar{m}$ forces $2(1-\ell) \geq (2-\lambda)^2 \geq 1$, where the second inequality is the minimum of $(2-\lambda)^2$. Therefore, $\ell \leq \frac{1}{2}$.

Lemma 5.A.8. *On its feasible regime, we have that* $p_3^\nabla(\ell, m, \lambda) \in [0,1]$.

Proof. We express $p_3^\nabla = \frac{N_3}{D_3}$, where the numerator and denominator are defined as follows:

$$N_3 := \lambda[2(1+\ell) - \lambda] - \frac{4(\ell - m)}{1-m};$$

$$D_3 := (1-\ell)(1+m)^2 - (\ell - m)(2-\lambda)^2.$$

Evaluating their gradients we get:

$$\frac{dD_3}{dm} = (2-\lambda)^2 + 2(1-\ell)(1+m);$$

$$\frac{dN_3}{dm} = \frac{4(1-\ell)}{(1-m)^2}.$$

We denote $\tilde{m} := \lambda - 1$. The lower end interval for $m \leq 0$ is $\max\{\bar{m}, \tilde{m}\}$.

1. **Proof of $p_3^\nabla \geq 0$.** We show that both denominator and numerator are nonnegative. Inspecting their expressions, we observe that they increase in m and therefore achieve their minimum at $\max\{\tilde{m}, \bar{m}\}$. At these points, one can check the following identities:

$$D_3(\tilde{m}) = (\tilde{m} - \bar{m})[(1-\lambda)^2 + 2\ell(1-\lambda) + 1];$$

$$D_3(\bar{m}) = -(\tilde{m} - \bar{m})\frac{(1-\ell)(2-\lambda)^2[2(1+\ell) - \lambda]}{(1-\lambda)^2 + 2\ell(1-\lambda)+1};$$

$$N_3(\tilde{m}) = \frac{D_3(\tilde{m})}{2-\lambda};$$

$$N_3(\bar{m}) = \frac{-D_3(\tilde{m})}{\lambda}.$$

The maximum between $\tilde{m}$ and $\bar{m}$ dictates the signs. For $\tilde{m} \geq \bar{m}$, it results that $D_3(\tilde{m}) \geq 0$, thus also $N_3(\tilde{m}) \geq 0$. Similarly, for $\tilde{m} \leq \bar{m}$, we get $D_3(\bar{m}) \geq 0$ and $D_3(\tilde{m}) \leq 0$, hence $N_3(\bar{m}) \geq 0$. In conclusion, both D_3 and N_3 are nonnegative.

2. **Proof of $p_3^\nabla \leq 1$.** We show $p_3^\nabla(\ell, m, \lambda) \leq p_3^\nabla(\ell, 0, \lambda) \leq 1$ by (i) demonstrating that p_3^∇ is increasing in m; and (ii) evaluating its maximum at $m = 0$.

i) We prove $\frac{dp_3}{dm} = \frac{N'(m)D(m)-N(m)D'(m)}{D^2(m)} \geq 0$ by showing that $H(m) \geq 0$, where $H(m) := (1-m^2)[N'(m)D(m) - N(m)D'(m)]$. We demonstrate its monotone increase with m and that it reaches its maximum for $m = 0$. Its first derivative $H'(m)$ reads as:

$$H'(m) = 2H_1(m)H_2(m),$$

where

$$H_1(m) = (2-\lambda)^2 + (1+3m)(1-\ell);$$

$$H_2(m) = (2-\lambda)(\lambda - 2\ell) + m[2(1-\ell)\lambda + (2-\lambda)^2].$$

The slope of $H_1(m)$ is nonnegative, hence

$$H_1(m) \geq H_1(\tilde{m}) \geq \big(\lambda - \frac{1}{2})^2 + \frac{5}{4} + (\frac{1}{2} - \ell) + 3\ell(1-\lambda) \geq 0.$$

On the other hand, $H_2(m)$ has a nonnegative slope, thus its minimum on the feasible range is reached at the lower bound $\max\{\tilde{m}, \bar{m}\}$, where

$$H_2(\tilde{m}) = (\tilde{m} - \bar{m})[(1-\lambda)^2 + 2\ell(1-\lambda) + 1];$$

$$H_2(\bar{m}) = -(\tilde{m} - \bar{m})2(1-\ell).$$

Hence, $H_2(\max\{\tilde{m}, \bar{m}\}) \geq 0$ and therefore $H(m) \geq 0$. Further on, evaluating H at $\max\{\tilde{m}, \bar{m}\}$ gives

$$H(\tilde{m}) = H_2(\tilde{m})^2 \geq 0;$$

$$H(\bar{m}) = \frac{8(1-\ell)^2(2-\lambda)(1+\ell-\lambda)}{[(1-\lambda)^2 + 2\ell(1-\lambda) + 1]^3} H_2(\tilde{m})^2 \geq 0.$$

Therefore, $H(m) \geq H(\max\{\tilde{m}, \bar{m}\}) \geq 0$ and hence p_3^∇ is increasing in m.

ii) Its maximum is obtained at $m = 0$, namely

$$p_3^\nabla(\ell, 0, \lambda) = 1 - \frac{(1-\ell)(1-\lambda)^2}{(1-\ell) - \ell(2-\lambda)^2}.$$

Due to $\lambda \leq 1$ and $\ell \leq \frac{1}{2}$, the denominator is nonnegative and hence $p_3^\nabla(\ell, 0, \lambda) \leq 1$. Nevertheless, substituting $m = 0$, $\lambda = 1$ and $\ell = \frac{1}{2}$ yields the maximum value $p_3^\nabla(\frac{1}{2}, 0, 1) = 1$.

□

Lemma 5.A.9. *The slack matrix A, defined in (5.20), is positive semidefinite with the choice $p = p_3^{\nabla}$ and non-zero multipliers $\alpha_{u_0,v_0} = 1$, $\alpha_{u_1,v_0} = 1 + \alpha$ and $\alpha_{v_0,u_1} = 2 + \alpha$, where $\alpha = \frac{-2m}{1+m}$.*

Proof. By construction, $\det A(\alpha) = 0$. After substitutions, the second principal minor becomes

$$\det A_{(2:3,2:3)} = \frac{[\ell(1-m) + 3 + m]\,[\lambda(1 - 2\ell m + m^2) - 2(\ell - m)(1-m)]^2}{\lambda^2(1-m)(1+m)^2 D_3},$$

where $D_3 \coloneqq (1-\ell)(1+m)^2 - (\ell - m)(2-\lambda)^2$ is proved to be nonnegative in the proof of Lemma 5.A.8. Consequently, $\det A_{(2:3,2:3)} \geq 0$.

Nevertheless, proving $A_{(3,3)} = 2(2+\alpha) - (\ell - m)p \geq 0$ reduces to checking that $p_3^{\nabla} \leq \frac{4}{(1+m)(\ell-m)}$. The latter results by bounding $\frac{4}{(1+m)(\ell-m)} \geq \frac{4}{1-m^2} \geq 4$ and employing $p_3^{\nabla} \leq 1$ (see Lemma 5.A.8). □

Chapter 6

Outline and future work

6.1 Summary

In this thesis, we provide tight analyses of standard first-order methods for smooth and composite optimization problems. The key methodology supporting these results is the performance estimation problem (PEP) framework, which both guided and confirmed our analyses. Beyond obtaining non-improvable convergence rates, an important contribution of this work is the development of systematic analytical proofs that deepen the understanding of worst-case scenarios.

In Chapter 3 we completely analyse **gradient descent**, the classical method for unconstrained optimization, for any constant stepsize choice, covering both nonconvex and convex functions. Our analysis is among the few to establish tight bounds for arbitrary stepsizes and curvatures, as it moves beyond simple one-step inequalities. Typically, due to their inherent difficulty for any stepsize choice, exact analytical studies are restricted to specific stepsize sequences, which streamline the proofs and usually provide better worst-case guarantees. Our proof technique is a sum-of-squares method, where the terms are constructed to capture the worst-case behaviour (as their cancellation provides necessary conditions for worst-case functions). The interpolation inequalities are linearly combined through certain multipliers, whose expressions were obtained by solving a linear system via a systematic cancellation procedure, in contrast to the common PEP-based approach of guessing them. Leveraging this analysis, we also propose a universal stepsize sequence that is independent of the iteration count. Preliminary worst-case experiments suggest this sequence is optimal for performance measures based on consecutive iterations, which is the only setting that yields iteration-count-independent guarantees so far. This sequence still has to be tested on relevant practical applications.

In Chapter 4 we apply the same complete analysis to the **proximal gradient descent (PGD)** algorithm, the natural extension to the constrained case, examining it through the lens of the **difference-of-convex algorithm (DCA)**. We establish tight analytical convergence rates for all curvature combinations of the two component functions; some rates are proven, while others are conjectured. The analysis is more challenging than that for gradient descent, as most rates required more complicated sum-of-squares decompositions. A key element in this construction was the inclusion of terms reflecting (sub)gradient differences and consecutive iterate distances. The conjectured expressions extend those for unconstrained gradient descent, and their demonstration is left as further work. Moreover, for standard convex constraints, we recover the same optimized constant and scheduled stepsize sequences as for gradient descent, resulting from identical rates and the same underlying descent lemma. Additionally, we show that PGD stepsize sequences can be translated to DCA settings where curvature shifting is beneficial (i.e., when the subtracted function is smooth). This allows for a DCA version with adaptive curvatures that vary at each iteration, contingent on the strong condition that the first function has an available exact convex conjugate. Practical validation of this adaptive version is left as future work.

In Chapter 5, for the **Douglas–Rachford splitting** method, we establish sublinear tight convergence rates under relaxed curvature assumptions, requiring only prox-boundedness of the nonconvex function, the most general smooth setting, and covering all admissible choices of stepsizes and relaxation parameters. The analysis, based on two standard performance measures (iterate progress and residual subgradient), improves upon existing state-of-the-art bounds.

In Table 6.1.1, we present a concise summary of the previously open problems that have been resolved, compare with the state of the art, and list the problems that remain open.

Table 6.1.1: State of the art prior to this thesis, thesis improvements, and open problems (compact summary).

Topic	Prior SoA & gaps in tight convergence analysis	This thesis	Open problems
GD	• Nonconvex: • μ not exploited • non-maximal stepsizes ($\leq \sqrt{3}/L$ vs $2/L$) [2] • no tight worst-cases for $\gamma L \in (1,2)$ • no proven worst-case-optimal stepsize (for $\mu \neq -L$) • Strongly convex: • no full stepsize-range proof (gradient norm vs initial objective gap)	• Complete constant stepsize analysis (weakly/strongly convex) in $\min_i \lVert\nabla f(x_i)\rVert^2$ (Theorems 3.2.1 to 3.2.3) • Tight worst-case constructions: 2D (Proposition 3.6.6), 3D (Proposition 3.6.8) (3D new in PEP) • Worst-case-optimal constant stepsizes (Proposition 3.3.3) • Horizon-free schedule (Theorem 3.3.1)	• Explicit worst-case functions (beyond interpolating triplets) for 2D/3D. • Empirical evaluation of parameter choices (especially horizon-free).
PGD/ DCA	• PGD: • weakly convex: limited analysis; typically $\gamma \leq 1/L$ [1] • strongly convex: existing analyses in other measures • DCA (convex–convex): • tight rates in our measures (regimes p_{1-2}, r_{1-2}) [3] • upper bound proofs technical; lower bounds partial.	• DCA extension: weak convexity in subtracted term $\Rightarrow$ PGD view with $\gamma > 1/L$. • PGD/DCA rates for (strongly) convex cases in nonstandard measure. • One-step DCA $\Rightarrow$ 6 regimes (distance-1 proofs): Theorem 4.2.1 (residual gradient) & Theorem 4.3.1 (gradient mapping). • Conjectured stepsize/curvature landscape ($\geq$1 smooth term); numerically supported: • residual gradient: Conjectures 4.2.3 to 4.2.5; partial proofs in Theorems 4.2.2 and 4.2.3 (strongly convex objectives); • gradient mapping: Conjecture 4.3.1; partial proofs in Theorems 4.3.2 and 4.3.3 (strongly convex objectives). • Curvature shifting for worst-case-optimal selection: Section 4.7	• Matching lower bounds for 4/6 regimes (DCA), both measures. • Distance-2 proofs for DCA; resolve conjectured domain: • DCA: $-\mu_2 < L_2 < \dfrac{-\mu_1\mu_2}{\mu_1+\mu_2}$. • PGD: γL_φ between first/last stepsize thresholds. • Empirical evaluation of curvature shifting. • Adapt/evaluate optimized PGD schedules for DCA (including “silver” stepsizes).

Continued on next page

Topic	Prior SoA & gaps in tight convergence analysis	This thesis	Open problems
DRS	• Smooth nonconvex + prox-bounded setting. • Convergence via DRE [76, 120]. • Tight parameter ranges (stepsize & relaxation); • non-tight constants [120].	• Exact worst-case constants (same SoA framework): • iterate progress: Theorem 5.3.2 • residual subgradient: Theorem 5.4.1 • Worst-case-optimal parameters (stepsize & relaxation) from iterate progress (Proposition 5.3.1 and Conjecture 5.3.1)	• Empirical evaluation of parameter choices. • Tight convex-case bounds for functions (operators known), without DRE. • DRE as Lyapunov: best measure? (extra residual regime).

6.2 PEPterns: Concluding remarks on PEP-based proofs

This section summarizes my personal insights on how to derive PEP-based proofs. The remarks are tailored to the performance criteria emphasized in this thesis, namely bounds on a (best) gradient norm under a bounded initial optimality gap (or suitable surrogates). In this setting, descent-type inequalities often yield performance guarantees through simple telescoping arguments, and the key quantities are typically easier to track than for other criteria, such as bounding $f_N - f_\star$ under an initial distance condition $\|x_0 - x_\star\| \le R$.

Summary of takeaways. The PEP viewpoint suggests a useful unifying template for deriving *exact* worst-case bounds for fixed-step first-order methods. In the end, the target statement is typically a quadratic inequality of the form

$$x^\top A x \;\ge\; 0,$$

where x stacks the relevant algorithmic quantities (iterates, gradients, subgradients, or auxiliary points) and A is a symmetric *slack matrix* obtained from a linear combination of interpolation inequalities. For instance, in our gradient descent analysis for best gradient norm subject to initial objective gap, one can take $x = [g_0,\ g_1,\ \ldots,\ g_N]^\top$, and feasibility of a proposed dual certificate reduces to showing $A \succeq 0$. This perspective clarifies why many exact analyses (including the ones in this thesis) naturally reduce to (i) identifying a small set of active inequalities and (ii) proving positive semidefiniteness of a structured matrix. This slack matrix structure is used in particular to derive the rates for DRS in Chapter 5, where an extended blueprint is provided in Section 5.5.3. However, it can also be used to reformulate the proofs for GD (Chapter 3), PGD/DCA (Chapter 4), and may be useful, as further work, to fill in the in PGD/DCA exact analysis, namely to prove the series of conjectures in Chapter 4 (Conjectures 4.2.4, 4.2.5, 4.3.1 and 4.6.2 for DCA or Conjectures 4.2.3, 4.6.1 and 4.6.3 for PGD).

Remarks on how the proofs are found. A recurring practical issue is that the dual multipliers are rarely unique. Consequently, a successful analytical derivation usually starts with a careful reduction of the interpolation system: remove redundant inequalities, keep a minimal set of *active* constraints, and use them as the guiding structure for the proof. In many cases, numerical PEP runs for small horizons ($N = 1, 2, 3, \ldots$) are sufficient to reveal which inequalities are active and to suggest a pattern for general N. This "pattern mining" is effective, but it must be used with caution: as the Gram matrix dimension grows, numerical accuracy deteriorates and small residuals become difficult to interpret (e.g., distinguishing a truly zero entry from a numerical artifact). This limitation disappears in fully symbolic derivations.

A second recurring phenomenon is that multi-step guarantees often follow from one-step identities. Once a suitable one-step inequality is established, the N-step statement is frequently obtained by telescoping (possibly after reweighting terms). This is particularly common when the performance measure is a sum of per-iteration decreases, or when the proof tracks a potential function.

A constructive route to worst-case values. A pragmatic way to guess the exact worst-case value is to assume that, at the worst case, the active interpolation inequalities hold with equality (preferably for a minimal active set). One can then algebraically "massage" these equalities until the performance metric appears explicitly, and complete squares with nonnegative coefficients. Discarding the remaining nonnegative terms (typically sums of squares with nonnegative coefficients), yields a candidate worst-case expression. When this procedure succeeds, it provides both (i) the candidate tight constant and (ii) a clear certificate structure explaining why the bound is tight.

General proof template. All PEP-based proofs used in this thesis follow the same logical structure: a nonnegative linear combination of interpolation inequalities yields an inequality of the form

$$\begin{aligned}&\text{(initial condition)} \;-\; \text{(performance metric)}\\ &\qquad\qquad -\; \sum_{i,j} \alpha_{i,j}\,\text{(interpolation inequality)} \;\geq\; 0,\end{aligned}$$

where $\alpha_{i,j}$ are multipliers of inequalities connecting iterations i and j. Embedding the initial condition (and any optional terms, such as an optimal value $f^\star$) into this combination typically produces

1. several explicit linear constraints on the multipliers $(\alpha_{i,j})$, and

2. a quadratic form representation $x^\top A(wc, \alpha)\, x \geq 0$, where the candidate worst-case value wc is embedded in the slack matrix A and α denotes the multipliers.

Establishing the bound then reduces to selecting (wc, α) such that $A(\alpha, wc) \succeq 0$. When an analytical expression for wc is available (e.g., from inspection of small-horizon numerics and an informed guess), the remaining task is to certify

$A \succeq 0$. Otherwise, wc can be obtained as the largest value for which $A(wc, \alpha)$ remains positive semidefinite for some admissible multipliers, as discussed in Remark 5.5.1; this viewpoint was useful in deriving the regimes $p_3^{\Delta,\nabla}$ for DRS in Chapter 5.

For compactness, the algebra is often carried out by treating the iterates and gradients as scalars. This does not restrict generality: the resulting inequalities depend only on inner products and can be lifted to arbitrary dimension through the Gram matrix representation.

Encoding worst-case structure. Exact bounds typically encode a specific worst-case structure (e.g., a quadratic model or a particular operator configuration). A principled approach is to identify the relationships characterizing the conjectured worst case and to enforce them in the certificate, either by completing squares or by an equivalent quadratic form representation. For example, if a regime is attained by the quadratic function $f(x) = \frac{L}{2}\|x\|^2$ (as suggested by small-horizon PEP experiments), then for gradient descent with stepsize γ the gradients satisfy $g_{k+1} = (1-\gamma L)g_k$. Consequently, cross terms such as $\langle g_{k+1}, g_k\rangle$ can be handled by introducing the residual square

$$\|g_{k+1} - (1-\gamma L)g_k\|^2 \;\geq\; 0,$$

which is essentially a Young-type completion; it makes the worst-case relation explicit and enables systematic cancellations.

Slack matrix manipulations. A crucial step is to rewrite the slack either (i) as a sum of squares (closely aligned with classical "Lyapunov"-style arguments) or (ii) as a quadratic form $x^\top Ax$ with $A \succeq 0$. The sum of squares representation is particularly informative: it often reveals the worst-case scenario by indicating which residuals must vanish at equality.

In some analyses (notably for DRS in Chapter 5), it is advantageous to encode worst-case relations *directly* at the level of the slack matrix, without constructing a sum of squares decomposition upfront. Concretely, one applies a linear change of variables $x \mapsto \bar{x} = Tx$ so that $\bar{x}$ isolates residual coordinates that vanish under the conjectured worst-case identities (see Section 5.5). Since these residuals appear as squared terms with nonnegative multipliers, they can be treated as discardable in the certificate; the remaining coordinates must then cancel through the choice of multipliers. The transformed matrix $T^\top AT$ is typically enforced to be singular and structured: canceling selected entries yields equations for both the multipliers and the worst-case value, while nonnegativity of the remaining principal minors delineates the parameter range of the active regime.

To illustrate the mechanism on a simple example, consider $x = [g_k,\ g_{k+1}]^\top$ and the relation $g_{k+1} = (1-\gamma L)g_k$. Define

$$\bar{x} = \begin{bmatrix} g_{k+1} \\ (1-\gamma L)g_k \end{bmatrix} = Tx, \qquad T = \begin{bmatrix} 1 & 0 \\ \frac{1}{1-\gamma L} & \frac{-1}{1-\gamma L} \end{bmatrix}.$$

If the original slack is $x^\top A(wc,\alpha)\,x$, then after this change of variables the slack matrix becomes

$$T^\top A T = \frac{1}{(1-\gamma L)^2}\begin{bmatrix} \begin{array}{l} A_{(2,2)} + 2(1-\gamma L)A_{(1,1)} \\ \quad + (1-\gamma L)^2 A_{(1,2)} \end{array} & -(1-\gamma L)A_{(1,2)} - A_{(2,2)} \\ \cdot & A_{(2,2)} \end{bmatrix}.$$

In this representation, the enforced worst-case identity yields a *singular* structure: positive semidefiniteness is governed by the remaining principal block $A_{(2,2)}$, while canceling the $(1,1)$ and $(1,2)$ entries yields equations determining the multipliers and the worst-case value. Moreover, the condition $A_{(2,2)} \geq 0$ delineates the parameter range in which the corresponding regime is active. This "transform-and-cancel" approach extends to multiple iterations and often reveals stable patterns in both the multiplier structure and the resulting worst-case expression.

Limitations and drawbacks. Despite their exactness, PEP-style proofs have several drawbacks:

1. **Non-uniqueness of multipliers.** Even for a minimal active set, dual multipliers are typically not unique, which can make symbolic identification difficult. Moreover, within a given regime, several subdomains may admit proofs with the same minimal number of *active* interpolation inequalities, while relying on different active sets. At regimes boundaries, however, the multipliers are unique because the adjacent proofs coincide.

2. **PSD verification can be nontrivial.** After cancellations, one still needs to prove $A \succeq 0$. The remaining matrix can be structured, but algebraically delicate, as observed in Section 5.5 for DRS.

3. **Low interpretability in the raw form.** Without additional rewriting (e.g., sum-of-squares), the resulting certificates can look far from standard descent-lemma analyses.

4. **Non-uniqueness of worst-case instances in higher dimension.** In one dimension, extremizers are often essentially unique up to trivial transformations,

while in higher dimension multiple non-equivalent worst-case constructions may coexist.

Possible extensions. These observations suggest several extensions and methodological directions:

- **Systematic "active-set" identification.** Developing principled rules to reduce interpolation systems and predict active inequalities could significantly shorten analytical derivations.
- **Automated symbolic PSD proofs for structured slacks.** Many slack matrices encountered here have low-rank or block structure; exploiting such structure could enable more routine certification.
- **Low-dimensional extremizers and rank-promoting heuristics.** When searching for worst-case instances, trace/log-det heuristics can sometimes promote low-rank Gram matrices; understanding when these heuristics succeed (or fail) remains an open practical question.
- **Regime boundaries and optimal tuning.** Empirically, optimized parameters often lie on boundaries between regimes, which may explain why tuning problems frequently reduce to checking finitely many candidate patterns. Formalizing this phenomenon could lead to simpler tuning rules.

Overall, the PEP framework is best viewed not only as an SDP that produces numerical worst-case certificates, but also as a proof design language: it suggests what to combine, what must cancel, and which algebraic identities encode the worst-case structure. The challenge is to translate these certificates into short, interpretable arguments that match the style and intuition of the optimization community.

6.3 Future work in tight convergence analysis

GD stepsizes. While we proposed a stepsize sequence that is independent of the iteration horizon N and appears optimal for guarantees based on consecutive inequalities, this sequence remains dependent on the chosen performance measure. For example, in the (strongly) convex case, tracking a complementary rate (such as ${f(x_N) - f_*}/{\|x_0 - x^*\|^2}$) requires reversing the sequence. This is consistent with the H-duality phenomenon [67], which is obscured when using constant stepsizes. This limitation highlights several key research directions. First, a primary goal is to find an optimized sequence of stepsizes that is independent of both the horizon N and the chosen performance measure. Second, such a "universal" schedule could be directly translated to Proximal Gradient Descent (PGD) settings, since, as we showed (with formalization left as future work), the descent lemmas from gradient descent can be translated to proximal gradient descent. Moreover, this dependency on performance measure raises a more fundamental question: what is the "best" (if one exists) performance criterion to study that provides the most meaningful guarantee of practical performance?

Additional methods. Leveraging new advancements in the interpolation inequalities [20], PEP formulations can include linear operators. This enables the exploration of additional primal-dual methods such as ADMM [53, 50, 22], Chambolle-Pock [29], and Condat-Vũ [31, 128]. For strongly convex cases, such tight results are partially available based on this PEP methodology [21, 20, 137], and have also been obtained for ADMM using the PEP-alternative integral quadratic constraint (IQC) framework [91].

AdaGrad-type methods. They use adaptive stepsizes computed from past gradient information and have become very popular. This adaptivity, however, poses a challenge for the standard Performance Estimation Problem (PEP) framework, which is designed for pre-scheduled, trajectory-independent stepsizes. Extending the PEP methodology to accommodate such adaptive stepsize selection remains a key research direction. Related work in this vein includes PEP analyses for Polyak stepsizes, with and without momentum [10] (notably requiring access to the optimal value), nonlinear conjugate gradient methods [33], and Barzilai-Borwein methods [12] on quadratic functions [19, §7.3].

Systematic analytical proofs. Current analyses of tight performance bounds are highly tailored. Each setup, defined by a specific performance measure, initial criterion, function class, and method, often requires a unique analytical approach. Although connections between different algorithms and bounds exist, as also observed in this work, they typically emerge only after extensive, separate analyses. A significant bottleneck in this process is inferring the analytical formulas for multipliers, worst-case values, and primal solutions. While symbolic regression methods [69] can aid this step, they demand highly accurate numerical solutions. Obtaining this accuracy is difficult, as the dimensionality of the underlying SDPs increases quadratically with the iteration count due to interpolation constraints. Therefore, a valuable research direction is the development of an automated method to obtain symbolic proofs for given optimization setups.

Bibliography

[1] Hadi Abbaszadehpeivasti. *Performance analysis of optimization methods for machine learning*. CentER dissertation. CentER, Tilburg University, 2024.

[2] Hadi Abbaszadehpeivasti, Etienne de Klerk, and Moslem Zamani. The exact worst-case convergence rate of the gradient method with fixed step lengths for L-smooth functions. *Optimization Letters*, 16(6):1649–1661, 2022.

[3] Hadi Abbaszadehpeivasti, Etienne de Klerk, and Moslem Zamani. On the rate of convergence of the difference-of-convex algorithm (dca). *Journal of Optimization Theory and Applications*, 202(1):475–496, Jul 2024.

[4] Hadi Abbaszadehpeivasti and Moslem Zamani. On the convergence rate of the Douglas-Rachford splitting algorithm. *arXiv preprint arXiv:2509.06676*, 2025.

[5] Jason Altschuler and Pablo Parrilo. Acceleration by stepsize hedging: Multi-step descent and the silver stepsize schedule. *J. ACM*, 72(2), 03 2025.

[6] Jason M. Altschuler and Pablo A. Parrilo. Acceleration by stepsize hedging: Silver stepsize schedule for smooth convex optimization. *Mathematical Programming*, 213(1):1105–1118, Sep 2025.

[7] Mohammad Askarizadeh, Alireza Morsali, Sadegh Tofigh, and Kim Khoa Nguyen. Convex-concave programming: An effective alternative for optimizing shallow neural networks. *IEEE Transactions on Emerging Topics in Computational Intelligence*, pages 1–14, 2024.

[8] Felipe Atenas, Claudia Sagastizábal, Paulo J. S. Silva, and Mikhail Solodov. A unified analysis of descent sequences in weakly convex optimization, including convergence rates for bundle methods. *SIAM Journal on Optimization*, 33(1):89–115, 2023.

[9] Pranjal Awasthi, Anqi Mao, Mehryar Mohri, and Yutao Zhong. Dc-programming for neural network optimizations. *Journal of Global Optimization*, Jan 2024.

[10] Mathieu Barré, Adrien Taylor, and Alexandre d'Aspremont. Complexity guarantees for polyak steps with momentum. In Jacob Abernethy and Shivani Agarwal, editors, *Proceedings of Thirty Third Conference on Learning Theory*, volume 125 of *Proceedings of Machine Learning Research*, pages 452–478. PMLR, 09–12 Jul 2020.

[11] Mathieu Barré, Adrien Taylor, and Francis Bach. Principled analyses and design of first-order methods with inexact proximal operators. *Mathematical Programming*, 201(1):185–230, 2022.

[12] Jonathan Barzilai and Jonathan M. Borwein. Two-point step size gradient methods. *IMA Journal of Numerical Analysis*, 8(1):141–148, 1 1988.

[13] Heinz Bauschke and Patrick Combettes. *Convex Analysis and Monotone Operator Theory in Hilbert Spaces.* Springer Publishing Company, Incorporated, 1st edition, 2011.

[14] Heinz Bauschke, Walaa Moursi, and Xianfu Wang. Generalized monotone operators and their averaged resolvents. *Mathematical Programming*, 189(1):55–74, Sep 2021.

[15] Amir Beck. *First-order methods in optimization.* SIAM-Society for Industrial and Applied Mathematics, 2017.

[16] Dimitri Bertsekas. *Convex optimization algorithms.* Athena Scientific, Belmont, MA, 2015.

[17] Dimitri Bertsekas. *Nonlinear programming.* Athena scientific optimization and computation series. Athena Scientific, 2016.

[18] Jérôme Bolte, Shoham Sabach, and Marc Teboulle. Proximal alternating linearized minimization for nonconvex and nonsmooth problems. *Mathematical Programming*, 146(1–2):459–494, 2014.

[19] Nizar Bousselmi. *Performance analysis of primal-dual and second-order optimization methods.* PhD thesis, Université catholique de Louvain, Louvain-la-Neuve, Belgium, 2025.

[20] Nizar Bousselmi, Julien Hendrickx, and François Glineur. Interpolation conditions for linear operators and applications to performance estimation problems. *SIAM Journal on Optimization*, 34(3):3033–3063, 2024.

[21] Nizar Bousselmi, Nelly Pustelnik, Julien Hendrickx, and François Glineur. Comparison of proximal first-order primal and primal-dual algorithms via performance estimation. In *2024 32nd European Signal Processing Conference (EUSIPCO)*, pages 2647–2651, 2024.

[22] Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, and Jonathan Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122, 2011.

[23] Stephen Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge University Press, 2004.

[24] Luis M. Briceño–Arias and Nelly Pustelnik. Theoretical and numerical comparison of first-order algorithms for cocoercive equations and smooth convex optimization. *Signal Processing*, 206:108900, 2023. Preprint arXiv:2101.06152 (2021).

[25] Sébastien Bubeck. Convex optimization: Algorithms and complexity. *Foundations and Trends in Machine Learning 8, 231-357*, 2015.

[26] Yair Carmon, John Duchi, Oliver Hinder, and Aaron Sidford. Lower bounds for finding stationary points I. *Mathematical Programming*, 184(1):71–120, 2020.

[27] Yair Carmon, John C. Duchi, Oliver Hinder, and Aaron Sidford. Accelerated methods for nonconvex optimization. *SIAM Journal on Optimization*, 28(2):1751–1772, 2018.

[28] Coralia Cartis, Nicholas Gould, and Philippe Toint. On the complexity of steepest descent, Newton's and regularized Newton's methods for nonconvex unconstrained optimization problems. *SIAM Journal on Optimization*, 20(6):2833–2852, 2010.

[29] Antonin Chambolle and Thomas Pock. A first-order primal-dual algorithm for convex problems with applications to imaging. *Journal of Mathematical Imaging and Vision*, 40(1):120–145, 2011.

[30] Patrick Combettes and Jean-Christophe Pesquet. Proximal splitting methods in signal processing. In Heinz H. Bauschke, Regina Burachik, Patrick L. Combettes, Veit Elser, D. Russell Luke, and Henry Wolkowicz, editors, *Fixed-Point Algorithms for Inverse Problems in Science and Engineering*, pages 185–212. Springer, 2011.

[31] Laurent Condat. A primal–dual splitting method for convex optimization involving Lipschitzian, proximable and linear composite terms. *SIAM Journal on Optimization*, 23(2):1063–1083, 2013.

[32] Shuvomoy Das Gupta. *Advances in computer-assisted design and analysis of first-order optimization methods and related problems.* PhD thesis, Massachusetts Institute of Technology, 2024.

[33] Shuvomoy Das Gupta, Robert M. Freund, Xu Andy Sun, and Adrien Taylor. Nonlinear conjugate gradient methods: worst-case convergence rates via computer-assisted analyses. *Mathematical Programming*, 213(1):1–49, Sep 2025.

[34] Shuvomoy Das Gupta, Bart Van Parys, and Ernest Ryu. Branch-and-bound performance estimation programming: a unified methodology for constructing optimal optimization methods. *Mathematical Programming*, 204(1):567–639, 2024.

[35] Ingrid Daubechies, Michel Defrise, and Christine De Mol. An iterative thresholding algorithm for linear inverse problems with a sparsity constraint. *Communications on Pure and Applied Mathematics*, 57(11):1413–1457, 2004.

[36] Damek Davis and Dmitriy Drusvyatskiy. Stochastic model-based minimization of weakly convex functions. *SIAM Journal on Optimization*, 29(1):207–239, 2019.

[37] Etienne de Klerk, François Glineur, and Adrien Taylor. On the worst-case complexity of the gradient method with exact line search for smooth strongly convex functions. *Optimization Letters*, 11(7):1185–1199, 2017.

[38] Welington de Oliveira. The ABC of DC programming. *Set-Valued and Variational Analysis*, 28(4):679–706, 2020.

[39] Tao Pham Dinh and Hoai An Le Thi. Convex analysis approach to d.c. programming: theory, algorithm and applications. In *Acta Mathematica Vietnamica*, volume 22, 1997.

[40] Jim Douglas and Henry Rachford. On the numerical solution of heat conduction problems in two and three space variables. *Transactions of the American Mathematical Society*, 82:421–439, 1956.

[41] Yoel Drori. *Contributions to the complexity analysis of optimization algorithms.* PhD thesis, Tel Aviv University, Tel Aviv, Israel, 2014.

[42] Yoel Drori. The exact information-based complexity of smooth convex minimization. *Journal of Complexity*, 39:1–16, 2017.

[43] Yoel Drori and Ohad Shamir. The complexity of finding stationary points with stochastic gradient descent. In *Proceedings of the 37th International*

Conference on Machine Learning, ICML 2020, volume 119 of *Proceedings of Machine Learning Research*, pages 2658–2667. PMLR, 2020.

[44] Yoel Drori and Marc Teboulle. Performance of first-order methods for smooth convex minimization: a novel approach. *Mathematical Programming*, 145(1-2):451–482, 2014.

[45] Dmitriy Drusvyatskiy. The proximal point method revisited. *arXiv preprint arXiv:1712.06038*, 2017.

[46] Jonathan Eckstein and Dimitri Bertsekas. On the Douglas–Rachford splitting method and the proximal point algorithm for maximal monotone operators. *Mathematical Programming*, 55(1–3):293–318, 1992.

[47] Brecht Evens, Pieter Pas, Puya Latafat, and Panagiotis Patrinos. Convergence of the preconditioned proximal point method and Douglas–Rachford splitting in the absence of monotonicity. *Mathematical Programming*, 214(1):247–301, Nov 2025.

[48] Oisin Faust, Hamza Fawzi, and James Saunderson. A Bregman divergence view on the difference-of-convex algorithm. In Francisco Ruiz, Jennifer Dy, and Jan-Willem van de Meent, editors, *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics*, volume 206 of *Proceedings of Machine Learning Research*, pages 3427–3439. PMLR, 25–27 Apr 2023.

[49] Marguerite Frank and Philip Wolfe. An algorithm for quadratic programming. *Naval Research Logistics Quarterly*, 3(1-2):95–110, 1956.

[50] Daniel Gabay and Bertrand Mercier. A dual algorithm for the solution of nonlinear variational problems via finite element approximation. *Computers & Mathematics with Applications*, 2(1):17–40, 1976.

[51] Saeed Ghadimi. Conditional gradient type methods for composite nonlinear and stochastic optimization. *Mathematical Programming*, 173(1):431–464, Jan 2019.

[52] Pontus Giselsson and Stephen Boyd. Linear convergence and metric selection for Douglas–Rachford splitting and ADMM. *IEEE Transactions on Automatic Control*, 62(2):532–544, 2016.

[53] Roland Glowinski and André Marroco. Sur l'approximation, par éléments finis d'ordre un, et la résolution, par pénalisation-dualité, d'une classe de problèmes de dirichlet non linéaires. *Revue Française d'Automatique, Informatique, Recherche Opérationnelle. Analyse Numérique*, 9(R2):41–76, 1975.

[54] Baptiste Goujaud. *Constructive approaches to worst-case complexity analyses of gradient methods for convex optimization: contributions, new insights, and novel results.* PhD thesis, Institut Polytechnique de Paris, 2024.

[55] Baptiste Goujaud, Aymeric Dieuleveut, and Adrien Taylor. On fundamental proof structures in first-order optimization (tutorial). In *Proceedings of the 62nd IEEE Conference on Decision and Control (CDC)*, pages 3023–3030. IEEE, 2023. Tutorial session.

[56] Baptiste Goujaud, Céline Moucer, François Glineur, Julien Hendrickx, Adrien Taylor, and Aymeric Dieuleveut. PEPit: computer-assisted worst-case analyses of first-order optimization methods in Python. *Mathematical Programming Computation*, 16(3):337–367, Sep 2024.

[57] Baptiste Goujaud, Céline Moucer, François Glineur, Julien Hendrickx, Adrien Taylor, and Aymeric Dieuleveut. PEPit: Documentation (version 0.4.0). https://pepit.readthedocs.io, 2024. Accessed: 2026-01-19.

[58] Benjamin Grimmer. Provably faster gradient descent via long steps. *SIAM Journal on Optimization*, 34(3):2588–2608, 2024.

[59] Benjamin Grimmer, Kevin Shu, and Alex Wang. Accelerated gradient descent via long steps. *arXiv preprint arXiv:2309.09961*, 2023.

[60] Benjamin Grimmer, Kevin Shu, and Alex Wang. Accelerated objective gap and gradient norm convergence for gradient descent via long steps. *INFORMS Journal on Optimization*, 7(2):156–169, 2025.

[61] Le Thi Hoai An, Le Hoai Minh, and Pham Dinh Tao. New and efficient DCA based algorithms for minimum sum-of-squares clustering. *Pattern Recognition*, 47(1):388–401, 2014.

[62] Reiner Horst and Nguyen Thoai. DC programming: overview. *Journal of Optimization Theory and Applications*, 103(1):1–43, October 1999.

[63] Michel Journée, Yurii Nesterov, Peter Richtárik, and Rodolphe Sepulchre. Generalized power method for sparse principal component analysis. *Journal of Machine Learning Research*, 11(15):517–553, 2010.

[64] Hamed Karimi, Julie Nutini, and Mark Schmidt. Linear convergence of gradient and proximal-gradient methods under the Polyak-Łojasiewicz condition. In Paolo Frasconi, Niels Landwehr, Giuseppe Manco, and Jilles Vreeken, editors, *Machine Learning and Knowledge Discovery in Databases*, pages 795–811, Cham, 2016. Springer International Publishing.

[65] Donghwan Kim and Jeffrey Fessler. Optimizing the efficiency of first-order methods for decreasing the gradient of smooth convex functions. *Journal of Optimization Theory and Applications*, 188(1):192–219, 2021.

[66] Donghwan Kim and Jeffrey A. Fessler. Optimized first-order methods for smooth convex minimization. *Mathematical Programming*, 159(1):81–107, 2016.

[67] Jaeyeon Kim, Asuman Ozdaglar, Chanwoo Park, and Ernest Ryu. Time-reversed dissipation induces duality between minimizing gradient norm and function value. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*, NIPS '23, Red Hook, NY, USA, 2024. Curran Associates Inc.

[68] Jungbin Kim. A proof of the exact convergence rate of gradient descent. *arXiv preprint arXiv:2412.04427*, 2025.

[69] Gabriel Kronberger, Bogdan Burlacu, Michael Kommenda, Stephen M. Winkler, and Michael Affenzeller. *Symbolic regression*. CRC Press, Boca Raton, FL, 2024.

[70] Delphine Lambert, Jean-Pierre Crouzeix, Hien Nguyen, and Jean-Jacques Strodiot. Finite convex integration. *Journal of Convex Analysis*, 11(1):131–146, 2004.

[71] Gert Lanckriet and Bharath Sriperumbudur. On the convergence of the concave-convex procedure. *Advances in Neural Information Processing Systems*, 22, 2009.

[72] Hoai An Le Thi, Van Ngai Huynh, and Tao Pham Dinh. Convergence analysis of difference-of-convex algorithm with subanalytic data. *Journal of Optimization Theory and Applications*, 179(1):103–126, Oct 2018.

[73] Hoai An Le Thi and Manh Cuong Nguyen. DCA based algorithms for feature selection in multi-class support vector machine. *Annals of Operations Research*, 249(1):273–300, Feb 2017.

[74] Hoai An Le Thi and Tao Pham Dinh. DC programming and DCA: thirty years of developments. *Mathematical Programming*, 169(1):5–68, May 2018.

[75] Laurent Lessard, Benjamin Recht, and Andrew Packard. Analysis and design of optimization algorithms via integral quadratic constraints. *SIAM Journal on Optimization*, 26(1):57–95, 2016.

[76] Guoyin Li and Ting Kei Pong. Douglas–Rachford splitting for nonconvex optimization with applications to nonconvex feasibility problems. *Mathematical Programming*, 159(1-2):371–401, 2016.

[77] Guoyin Li and Ting Kei Pong. Global convergence of splitting methods for nonconvex composite optimization. *SIAM Journal on Optimization*, 26(4):2130–2167, 2016.

[78] Pierre-Louis Lions and Bertrand Mercier. Splitting algorithms for the sum of two nonlinear operators. *SIAM Journal on Numerical Analysis*, 16(6):964–979, 1979.

[79] Thomas Lipp and Stephen Boyd. Variations and extension of the convex–concave procedure. *Optimization and Engineering*, 17(2):263–287, Jun 2016.

[80] Boris Mordukhovich. *Variational analysis and generalized differentiation I: basic theory*, volume 330 of *Grundlehren der Mathematischen Wissenschaften*. Springer, Berlin, 2006.

[81] Jean Jacques Moreau. Proximité et dualité dans un espace hilbertien. *Bulletin de la Société Mathématique de France*, 93:273–299, 1965.

[82] Walaa M. Moursi and Lieven Vandenberghe. Douglas–Rachford splitting for the sum of a Lipschitz continuous and a strongly monotone operator. *Journal of Optimization Theory and Applications*, 183(1):179–198, 2019.

[83] Arkadi Nemirovski. Information-based complexity of linear operator equations and convex programming problems. In P. Pardalos, editor, *Lecture Notes in Economics and Mathematical Systems*, volume 397, pages 1–23. Springer, Berlin, Heidelberg, 1992.

[84] Arkadi Nemirovski. Efficiency of algorithms for solving convex programming problems. In *Proceedings of the International Congress of Mathematicians (ICM 1994)*, volume 1, pages 1599–1620. Birkhäuser, Zürich, 1994.

[85] Arkadi Nemirovski and David Yudin. *Problem complexity and method efficiency in optimization*. John Wiley & Sons, New York, 1983. Translated from the Russian by E. R. Dawson.

[86] Yurii Nesterov. A method for solving the convex programming problem with convergence rate $O(1/k^2)$. *Soviet Mathematics Doklady*, 27(2):372–376, 1983.

[87] Yurii Nesterov. How to make the gradients small. *Optima 88*, 2012.

[88] Yurii Nesterov. Gradient methods for minimizing composite functions. *Mathematical Programming*, 140(1):125–161, Aug 2013.

[89] Yurii Nesterov. *Lectures on convex optimization.* Springer Publishing Company, Incorporated, 2nd edition, 2018.

[90] Edward Duc Hien Nguyen, Jaewook Suh, Xin Jiang, and Shiqian Ma. Exact worst-case convergence rates for Douglas–Rachford and Davis–Yin splitting methods. *arXiv preprint arXiv:2506.23475*, 2025.

[91] Robert Nishihara, Laurent Lessard, Ben Recht, Andrew Packard, and Michael Jordan. A general analysis of the convergence of ADMM. In Francis Bach and David Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 343–352, Lille, France, 07–09 Jul 2015. PMLR.

[92] Evgeni Alekseevich Nurminskii. The quasigradient method for the solving of the nonlinear programming problems. *Cybernetics*, 9(1):145–150, 1973.

[93] Neal Parikh and Stephen Boyd. Proximal algorithms. *Foundations and Trends in Optimization*, 1(3):127–239, 2014.

[94] Panagiotis Patrinos and Alberto Bemporad. Proximal Newton methods for convex composite optimization. In *52nd IEEE Conference on Decision and Control*, pages 2358–2363, 2013.

[95] Rene Poliquin and Tyrrell Rockafellar. Prox-regular functions in variational analysis. *Transactions of the American Mathematical Society*, 348(5):1805–1838, 1996.

[96] Tyrrell Rockafellar. *Convex analysis.* Princeton Mathematical Series. Princeton University Press, Princeton, N. J., 1970.

[97] Tyrrell Rockafellar. Monotone operators and the proximal point algorithm. *SIAM Journal on Control and Optimization*, 14(5):877–898, 1976.

[98] Tyrrell Rockafellar and Roger Wets. *Variational analysis*, volume 317 of *Grundlehren der Mathematischen Wissenschaften.* Springer, Berlin, 3rd printing edition, 2009.

[99] Teodor Rotaru, François Glineur, and Panagiotis Patrinos. Exact worst-case convergence rates of gradient descent: a complete analysis for all constant stepsizes over nonconvex and convex functions. *Mathematical Programming*, 2026. doi:10.1007/s10107-025-02313-1. Online first.

[100] Teodor Rotaru, François Glineur, and Panagiotis Patrinos. Exact worst-case convergence rates of gradient descent: a complete analysis for all constant stepsizes over nonconvex and convex functions. *arXiv preprint arXiv:2406.17506*, 2024.

[101] Teodor Rotaru, Panagiotis Patrinos, and François Glineur. Tight analysis of difference-of-convex algorithm (DCA) improves convergence rates for proximal gradient descent. In *The 28th International Conference on Artificial Intelligence and Statistics*, volume 258, pages 4114–4122, 2025.

[102] Teodor Rotaru, Panagiotis Patrinos, and François Glineur. Tight convergence rates in gradient mapping for the difference-of-convex algorithm. In Hoai An Le Thi, Tao Pham Dinh, and Hoai Minh Le, editors, *Modelling, Computation and Optimization in Information Systems and Management Sciences*, pages 61–73, Cham, 2026. Springer Nature Switzerland.

[103] Teodor Rotaru, Panagiotis Patrinos, and François Glineur. Improved convergence rates for the difference-of-convex algorithm. *arXiv preprint arXiv:2403.16864*, 2024.

[104] Ernest Ryu, Adrien Taylor, Carolina Bergeling, and Pontus Giselsson. Operator splitting performance estimation: Tight contraction factors and optimal parameter selection. *SIAM Journal on Optimization*, 30(3):2251–2271, 2020.

[105] Lorenzo Stella, Andreas Themelis, and Panagiotis Patrinos. Forward–backward quasi-Newton methods for nonsmooth optimization problems. *Computational Optimization and Applications*, 67(3):443–487, Jul 2017.

[106] Kaizhao Sun and Xu Andy Sun. Algorithms for difference-of-convex programs based on difference-of-Moreau-envelopes smoothing. *INFORMS Journal on Optimization*, 5(4):321–339, 2022.

[107] Youran Sun, Yihua Liu, and Yi-Shuai Niu. Understand the effectiveness of shortcuts through the lens of DCA. *arXiv preprint arXiv:2412.09853*, 2024.

[108] Ksenia Syrtseva, Welington de Oliveira, Sophie Demassey, and Wim van Ackooij. Minimizing the difference of convex and weakly convex functions via bundle method. *Pacific Journal of Optimization*, 20(4):699–741, 2024.

[109] Pham Dinh Tao and Le Thi Hoai An. A d.c. optimization algorithm for solving the trust-region subproblem. *SIAM Journal on Optimization*, 8(2):476–505, 1998.

[110] Adrien Taylor. *Convex interpolation and performance estimation of first-order methods for convex optimization.* PhD thesis, UCLouvain - Université catholique de Louvain, Louvain-la-Neuve, 2017.

[111] Adrien Taylor. *Towards principled and systematic approaches to the analysis and design of optimization algorithms.* PSL Research University, 2024. Habilitation à diriger des recherches.

[112] Adrien Taylor and Yoel Drori. An optimal gradient method for smooth strongly convex minimization. *Mathematical Programming*, pages 1–38, 2022.

[113] Adrien Taylor, Julien Hendrickx, and François Glineur. Exact worst-case performance of first-order methods for composite convex optimization. *SIAM Journal on Optimization*, 27(3):1283–1313, 2017.

[114] Adrien Taylor, Julien Hendrickx, and François Glineur. Exact worst-case convergence rates of the proximal gradient method for composite convex minimization. *Journal of Optimization Theory and Applications*, 178(2):455–476, Aug 2018.

[115] Adrien Taylor, Julien Hendrickx, and François Glineur. Performance estimation toolbox (PESTO): Automated worst-case analysis of first-order optimization methods. In *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, pages 1278–1283, 2017.

[116] Adrien Taylor, Julien Hendrickx, and François Glineur. Smooth strongly convex interpolation and exact worst-case performance of first-order methods. *Mathematical Programming*, 161(1-2):307–345, 2017.

[117] Adrien Taylor, Bryan Van Scoy, and Laurent Lessard. Lyapunov functions for first-order methods: Tight automated convergence guarantees. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*, pages 4897–4906. PMLR, 2018.

[118] Marc Teboulle and Yakov Vaisbourd. An elementary approach to tight worst case complexity analysis of gradient based methods. *Mathematical Programming*, 201(1–2):63–96, 2022.

[119] Andreas Themelis, Ben Hermans, and Panagiotis Patrinos. A new envelope function for nonsmooth DC optimization. In *2020 59th IEEE Conference on Decision and Control (CDC)*, pages 4697–4702, 2020.

[120] Andreas Themelis and Panagiotis Patrinos. Douglas–Rachford splitting and ADMM for nonconvex optimization: Tight convergence results. *SIAM Journal on Optimization*, 30(1):149–181, 2020.

[121] Andreas Themelis, Luca Stella, and Panagiotis Patrinos. Forward–backward envelope for the sum of two nonconvex functions: further properties and nonmonotone linesearch algorithms. *SIAM Journal on Optimization*, 28(3):2274–2303, 2018.

[122] John Toland. A duality principle for non-convex optimisation and the calculus of variations. *Archive for Rational Mechanics and Analysis*, 71(1):41–61, May 1979.

[123] Manu Upadhyaya, Sebastian Banert, Adrien Taylor, and Pontus Giselsson. Automated tight Lyapunov analysis for first-order methods. *Mathematical Programming*, 209(1):133–170, 2025.

[124] Manu Upadhyaya, Adrien Taylor, Sebastian Banert, and Pontus Giselsson. AutoLyap: A python package for computer-assisted Lyapunov analyses for first-order methods. *arXiv preprint arXiv:2506.24076*, 2025.

[125] Bryan Van Scoy, Randy Freeman, and Kevin Lynch. The fastest known globally convergent first-order method for minimizing strongly convex functions. *IEEE Control Systems Letters*, 2(1):49–54, 2018.

[126] Pierre Vernimmen and François Glineur. Worst-case convergence analysis of relatively inexact gradient descent on smooth convex functions. In *Proceedings of the 32nd European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning (ESANN 2024)*, Bruges, Belgium, 2024. i6doc.com.

[127] Xuan Thanh Vo, Hoai An Le Thi, Tao Pham Dinh, and Thi Bich Thuy Nguyen. DC programming and DCA for dictionary learning. In Manuel Núñez, Ngoc Thanh Nguyen, David Camacho, and Bogdan Trawiński, editors, *Computational Collective Intelligence*, pages 295–304, Cham, 2015. Springer International Publishing.

[128] Băng Công Vũ. A splitting algorithm for dual monotone inclusions involving cocoercive operators. *Advances in Computational Mathematics*, 38(3):667–681, 2013.

[129] Kuaini Wang, Wenxin Zhu, and Ping Zhong. Robust support vector regression with generalized loss function and applications. *Neural Processing Letters*, 41(1):89–106, Feb 2015.

[130] Hassler Whitney. Analytic extensions of differentiable functions defined in closed sets. *Hassler Whitney Collected Papers*, pages 228–254, 1934.

[131] Hassler Whitney. Differentiable functions defined in closed sets. I. *Transactions of the American Mathematical Society*, 36:369–387, 1934.

[132] Liming Yang and Yannan Qian. A sparse logistic regression framework by difference of convex functions programming. *Applied Intelligence*, 45(2):241–254, Sep 2016.

[133] Chaorui Yao and Xin Jiang. A globally convergent difference-of-convex algorithmic framework and application to log-determinant optimization problems. *arXiv preprint arXiv:2306.02001*, 2023.

[134] Penghang Yin, Yifei Lou, Qi He, and Jack Xin. Minimization of $\ell_1 - \ell_2$ for compressed sensing. *SIAM Journal on Scientific Computing*, 37(1):A536–A563, 2015.

[135] Alan L Yuille and Anand Rangarajan. The concave-convex procedure (CCCP). In T. Dietterich, S. Becker, and Z. Ghahramani, editors, *Advances in Neural Information Processing Systems*, volume 14. MIT Press, 2001.

[136] Alp Yurtsever and Suvrit Sra. CCCP is Frank-Wolfe in disguise. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 35352–35364. Curran Associates, Inc., 2022.

[137] Moslem Zamani, Hadi Abbaszadehpeivasti, and Etienne de Klerk. The exact worst-case convergence rate of the alternating direction method of multipliers. *Mathematical Programming*, 208(1-2):243–276, 2024.

[138] Moslem Zamani and François Glineur. Exact convergence rate of the last iterate in subgradient methods. *SIAM Journal on Optimization*, 35(3):2182–2201, 2025.

[139] Zehao Zhang and Rujun Jiang. Accelerated gradient descent by concatenation of stepsize schedules. *arXiv preprint arXiv:2410.12395*, 2024.

[140] Zihan Zhang, Jason Lee, Simon Du, and Yuxin Chen. Anytime acceleration of gradient descent. In Nika Haghtalab and Ankur Moitra, editors, *Proceedings of Thirty Eighth Conference on Learning Theory*, volume 291 of *Proceedings of Machine Learning Research*, pages 5991–6013. PMLR, 30 Jun–04 Jul 2025.

Statement on the use of Generative AI

I did not use generative AI assistance tools during the research/writing process of my thesis, except for mere language assistance.
The text, code, and images in this thesis are my own (unless otherwise specified). Generative AI has only been used in accordance with the KU Leuven guidelines and appropriate references have been added. I have reviewed and edited the content as needed and I take full responsibility for the content of the thesis.

Curriculum vitae

May 26, 1994	Born in Piatra Neamț, Romania
2013-2017	B.Sc. in Electrical Engineering - Control Systems University Politehnica of Bucharest, Romania GPA: 10/10
2017-2019	M.Sc. in Complex Systems University Politehnica of Bucharest, Romania GPA: 10/10
2018-2021	M.Sc. in Computational Science and Engineering Technical University of Munich, Germany GPA: 1.7/1.0
Since 2021	Ph.D. in Engineering Science, jointly at Department of Electrical Engineering (ESAT) KU Leuven, Belgium Institute of Information and Communication Technologies, Electronics, and Applied Mathematics UCLouvain, Belgium

Publications

Journal papers

[100] Teodor Rotaru, François Glineur, and Panagiotis Patrinos. Exact worst-case convergence rates of gradient descent: a complete analysis for all constant stepsizes over nonconvex and convex functions. *arXiv preprint arXiv:2406.17506*, 2024 (accepted for publication in the journal Mathematical Programming [99], doi:10.1007/s10107-025-02313-1)

Conference papers

[101] Teodor Rotaru, Panagiotis Patrinos, and François Glineur. Tight analysis of difference-of-convex algorithm (DCA) improves convergence rates for proximal gradient descent. In *The 28th International Conference on Artificial Intelligence and Statistics*, volume 258, pages 4114–4122, 2025

[102] Teodor Rotaru, Panagiotis Patrinos, and François Glineur. Tight convergence rates in gradient mapping for the difference-of-convex algorithm. In Hoai An Le Thi, Tao Pham Dinh, and Hoai Minh Le, editors, *Modelling, Computation and Optimization in Information Systems and Management Sciences*, pages 61–73, Cham, 2026. Springer Nature Switzerland

UCLouvain
INSTITUTE OF INFORMATION AND COMMUNICATION TECHNOLOGIES, ELECTRONICS AND APPLIED MATHEMATICS (ICTEAM)
DEPARTMENT OF MATHEMATICAL ENGINEERING (INMA)
Avenue Georges Lemaître 4 bte L4.05.01
B-1348 Louvain-la-Neuve
teodor.rotaru@uclouvain.be

KU Leuven
FACULTY OF ENGINEERING SCIENCE
DEPARTMENT OF ELECTRICAL ENGINEERING
Kasteelpark Arenberg 10 box 2440
B-3001 Leuven
teodor.rotaru@kuleuven.be