



Institute for Information
and Communication Technologies,
Electronics and Applied Mathematics

Interpolation conditions for exact computer-aided analysis of optimization methods

Anne Rubbens

Thesis submitted in partial fulfillment
of the requirements for the degree of
Docteur en sciences de l'ingénieur-e

Dissertation committee:

Prof. Julien M. Hendrickx (UCLouvain, advisor)

Prof. François Glineur (UCLouvain)

Prof. Alex Olshevsky (Boston University)

Prof. Silvia Villa (University of Genoa)

Dr. Yoel Drori (Google Inc.)

Prof. Pierre-Antoine Absil (UCLouvain, chair)

Version of September 2025.

This template is powered by L^AT_EX.
It was created by Pierre-Yves Gousenbourger (2020).
© Anne Rubbens, UCLouvain (2025), all rights reserved.

Abstract

This thesis considers the tight performance analysis of optimization methods. It first aims at better understanding interpolation conditions for classes of functions or operators, that is, conditions ensuring consistency of datasets with a function or operator in the class. Combined with computer-aided approaches as, e.g., the Performance Estimation Problem (PEP) framework, these allow for exact performance analysis of given methods on these classes. It then extends the PEP framework to enable the analysis of broader optimization settings.

The thesis is structured in three parts. The first part introduces an algebraic approach to function and operator interpolation, and uses it to refine existing descriptions of a few function and operator classes, including weakly convex functions, smooth functions satisfying a Lojasiewicz condition, blockwise smooth functions, and Lipschitz monotone operators.

The second part focuses on second-order optimization. It derives interpolation conditions for univariate function classes with second-order properties, such as self-concordance and Lipschitz continuity of the Hessian. These conditions are used to obtain exact worst-case performance bounds on second-order methods in the univariate setting.

Finally, the third part extends the PEP framework to the tight analysis of stochastic optimization methods. Given a stochastic first-order method, a function class, and a noise model with specified expectation and variance, it proposes a family of semidefinite formulations that yield increasingly strong convergence guarantees on the problem, at the expense of an increasing complexity. The largest formulations, whose size grow exponentially with the number of iterations, provide tight guarantees.

Acknowledgments

UN doctorat, c'est une aventure jalonnée de doutes et de questionnements, que je n'aurais pas pu mener à bien sans le soutien d'une farandole de personnes.

Tout d'abord, merci à mon promoteur, Julien, pour son soutien tout au long de ma thèse, son enthousiasme lors de chaque discussion scientifique, et sa disponibilité pour me rassurer sur mes résultats et tout autre sujet. Merci également à François, Seb et Nizar, avec qui j'ai eu la chance de collaborer, qui m'ont poussée à persévérer, et m'ont brillamment démontré que la recherche, ensemble, est beaucoup plus enrichissante et intéressante. Merci aussi à Adrien, pour m'avoir si chaleureusement accueillie lors de mon séjour de recherche, et pour toutes les discussions enthousiasmantes qui ont suivi, tant scientifiques qu'informelles. Ce fut un réel plaisir de travailler et d'échanger avec toi. I would also like to thank all the members of my committee, Silvia Villa, Yoel Drori, Alex Olshevsky, François and Pierre-Antoine, for taking the time to so carefully read my thesis, and for the kind and insightful feedback.

Je voudrais ensuite remercier toutes celles et ceux qui ont contribué à faire de mon environnement de travail, l'Euler, un espace si chaleureux. Merci à Marie-Christine, Pascale, Astrid, John et François, qui prennent tant soin de nous, et à Etienne, pour m'avoir fait prendre l'air. Merci à Charles, Martina et Jana, pour l'ambiance toujours pleine de vie et de potins du a.019. Merci à la Team JUL, Renato Charles Martina Tonio Seb Yassine Nizar Erwan, pour les retraites toujours pleines de rebondissements, les fous rires et l'entraide. Un merci tout particulier à Yassine, qui a su me rassurer avant chaque étape de ce doctorat, à Seb, pour les aventures à vélo, et à Nizar, pour toutes les discussions et les conseils. Merci aussi à tous-tes les autres habitant-es de l'Euler, et en particulier Clem, Julien,

Philou, Astrid, Thomas, Antoine, Simon, Téo, Rémi, Guillaume, Virginie, Adrien, Dâna, Bastien, qui, à grand renforts de pauses café, papote, ou sorties pioupious, ont fait souffler sur ces années de travail comme un vent de vacances.

Je remercie également toute la Horde des Ingés en Transition, pour l'énergie, la motivation et les moments partagés. S'interroger sur le sens de notre métier, rêver à d'autres possibles, et s'activer pour un changement avec vous sont tant de projets qui m'ont portée tout au long de ces quatre années. Simon, merci pour tous les trajets vélo papote!

Enfin, merci à ma famille et mes ami-es sans qui je ne serais pas arrivée au bout de l'aventure. En particulier, merci à mes colocos Ed, Seb et Erica, pour m'avoir nourrie, encouragée et divertie tout au long des mois d'écriture. Merci à Eileen et Lara, mes deux petits soleils. Et merci, Vic, pour le soutien indéfectible de jour comme de nuit.

Contents

| | |
|--|-----------|
| Contents | iv |
| List of symbols | vi |
| List of acronyms | x |
| 1 Introduction | 1 |
| 1.1 Outline and contributions | 8 |
| 2 Interpolation conditions for function and operator classes . . . | 11 |
| 2.1 Global characterizations of function classes | 12 |
| 2.2 Discrete characterizations of function classes | 14 |
| 2.3 Operator classes and their discrete characterizations | 17 |
| 2.4 State of the art in interpolation conditions | 19 |
| 2.5 Challenges in obtaining interpolation conditions | 21 |
| 3 Computer-aided computation of worst-case performance guarantees | 22 |
| 3.1 Key elements in Performance Estimation | 23 |
| 3.2 Problem classes encompassed by the PEP framework | 26 |
| 3.3 A potential functions approach to performance analysis | 27 |
| 3.4 Other applications to computer-aided analysis | 28 |
| 3.5 State of the art in computer-aided performance analysis | 28 |

I An algebraic approach to function and operator interpolation

| | | |
|----------|--|------------|
| 4 | A constraint-based approach to interpolation, and weakly convex optimization. | 32 |
| 4.1 | Codes. | 33 |
| 4.2 | A constraint-based approach to interpolation | 34 |
| 4.3 | Pointwise extensibility | 40 |
| 4.4 | Proving non-interpolability of a constraint | 44 |
| 4.5 | Weakly convex optimization | 48 |
| 4.6 | Exhaustive analysis of a class of constraints | 61 |
| 4.7 | Concluding observations | 66 |
| 5 | A constructive approach to interpolation conditions | 68 |
| 5.1 | Codes | 70 |
| 5.2 | Iterative strengthening of a constraint | 70 |
| 5.3 | One-point strengthenings in practice | 74 |
| 5.4 | Examples: smooth convex functions and Lipschitz operators | 74 |
| 5.5 | Applications: improved descriptions of problem classes | 80 |
| 5.6 | Concluding observations | 102 |
| | Appendix to Chapter 5 | 104 |
| 5.A | Proof of Proposition 5.5 | 104 |
| 5.B | Proof of Proposition 5.7 | 105 |
| 5.C | Proof of Proposition 5.9 | 105 |

II Interpolation conditions for second-order univariate optimization

| | | |
|----------|--|------------|
| 6 | Generic method for univariate interpolation | 108 |
| 6.1 | Order m interpolation conditions | 109 |
| 6.2 | Overview of the method | 110 |
| 6.3 | Interpolation without function values | 113 |
| 6.4 | Interpolation with function values (2 points) | 113 |
| 6.5 | Interpolation with function values (all points) | 117 |
| 6.6 | Main Theorem | 118 |
| 7 | Computer-aided analysis of second-order univariate optimization | 120 |
| 7.1 | Codes. | 121 |

| | | |
|--|--|------------|
| 7.2 | Motivation | 121 |
| 7.3 | Interpolation of generalized self-concordant functions | 124 |
| 7.4 | Performance estimation of second-order optimization | 147 |
| 7.5 | Concluding observations | 160 |
| Appendix to Chapter 7 | | 161 |
| 7.A | Proof of Theorems 7.4 and 7.5 | 161 |

III Interpolation conditions for stochastic first-order optimization

| | | |
|--|--|------------|
| 8 | Computer-aided analysis of stochastic first-order methods . . . | 172 |
| 8.1 | Motivation | 173 |
| 8.2 | Preliminaries: computer-aided analysis of inexact methods. . . | 180 |
| 8.3 | Framework for automated analysis of stochastic methods . . . | 183 |
| 8.4 | Applications: performance analysis of SGD on $\mathcal{F}_{\mu,L}$ | 205 |
| 8.5 | Concluding observations | 216 |
| Appendix to Chapter 8 | | 218 |
| 8.A | Practical implementation of $(P_{\text{relax},1})$ and $(P_{\text{relax},\text{sym}})$ | 218 |

IV Conclusion

| | | |
|-------------------------------|---|------------|
| 9 | Research outcomes and perspectives | 223 |
| 9.1 | Research outcomes | 224 |
| 9.2 | Perspectives and open questions | 227 |
| Bibliography | | 231 |

List of symbols

- $|\cdot|$ _____ Absolute value (p. 11)
- $\|\cdot\|$ _____ Standard norm (p. 11)
- $\langle \cdot, \cdot \rangle$ _____ Standard inner product (p. 11)
- $\mathbf{1}_d$ _____ Vector of ones in \mathbb{R}^d (p. 203)
- Id_d _____ Identity matrix in \mathbb{S}^d (p. 122)
- \succcurlyeq _____ Löwner ordering (p. 23)
- \emptyset _____ Empty set (p. 11)
- $\nabla f(x)$ _____ Gradient of function f at x (p. 2)
- $\nabla^2 f(x)$ _____ Hessian of function f at x (p. 2)
- $\partial f(x)$ _____ Subdifferential of function f at x (p. 12)
- $\{\mathcal{A}_k\}_{k=0,1,\dots,N}$ _____ Filtration of available information at iteration k (p. 173)
- $\mathcal{B}_{c,R}$ _____ Ball of center c and radius R (p. 51)
- C_{init} _____ Initial condition (p. 5)
- \mathcal{C}^m _____ Class of m times continuously differentiable functions (p. 109)
- C_{variance} _____ Variance conditions (p. 188)
- C_{\perp} _____ Uncorrelation conditions (p. 188)
- $\text{dom}f$ _____ Effective domain of function f (p. 11)

- \mathcal{D}^d _____ Cone of diagonal matrices in \mathbb{S}^d (p. 173)
- ε_k _____ Random variable (p. 4)
- \mathbf{e}_k _____ k^{th} unit vector (p. 23)
- $E_{\varepsilon_k}[\cdot]$ _____ Expectation with respect to ε_k (p. 4)
- \mathcal{F} _____ Function class (p. 5)
- $\mathcal{F}_{\mu,L}(\mathbb{R}^d)$ _____ Class of L -smooth μ -strongly convex on \mathbb{R}^d (p. 2)
- $\mathcal{F}_{\mu,L}(C)$ _____ Class of L -smooth μ -strongly convex on $C \subseteq \mathbb{R}^d$ (p. 13)
- $\mathcal{F}_p(\mathcal{X})$ _____ Class of mappings satisfying p on $\mathcal{X} \subseteq \mathbb{R}^d$ (p. 34)
- $\mathcal{F}_p^m(\mathcal{X})$ _____ Class of maximal mappings satisfying p on $\mathcal{X} \subseteq \mathbb{R}^d$ (p. 36)
- $\mathcal{F}_{p-\mu,B}^m(\mathbb{R}^d)$ _____ Class of weakly convex functions
with bounded subgradients on \mathbb{R}^d (p. 48)
- $\mathcal{F}_{\mathfrak{L},L}(\mathbb{R}^d)$ _____ Class of L -smooth functions satisfying
a quadratic Łojasiewicz condition on \mathbb{R}^d (p. 81)
- $\mathcal{F}_{0,\mathbb{L}}(\mathbb{R}^d)$ _____ Class of convex blockwise \mathbb{L} -smooth functions on \mathbb{R}^d (p. 87)
- $\mathcal{F}_{\mu,q}(\mathbb{R}^d)$ _____ Class of uniformly convex functions on \mathbb{R}^d (p. 97)
- \mathcal{F}_M _____ Class of univariate M -Lipschitz functions (p. 110)
- $\int \mathcal{F}_M$ _____ Class of univariate M -smooth functions (p. 110)
- $\int^{(2)} \mathcal{F}_{M,\alpha,(+)}$ _____ Class of univariate generalized
self-concordant functions (p. 125)
- Graph F _____ Graph of an operator F (p. 36)
- \mathcal{H}_M _____ Class of univariate functions with M -Lipschitz Hessian (p. 122)
- $\mathcal{H}_{M,0}$ _____ Class of univariate convex functions
with M -Lipschitz Hessian (p. 122)
- \mathcal{M} _____ Method (p. 5)
- \mathbb{N} _____ Natural numbers (p. 11)
- $[N]$ _____ $\{1, \dots, N\}$ (p. 11)

- \mathcal{P} _____ Performance measure (p. 5)
- p^{ij} _____ Evaluation of a pairwise constraint at (x_i, x_j) (p. 40)
- p^{xy} _____ Evaluation of a pairwise constraint at (x, y) (p. 40)
- $\tilde{P}_{K, \bar{\mathcal{X}}}$ _____ K -wise one-point strengthening of
constraint p with respect to $\bar{\mathcal{X}} \subseteq \mathbb{R}^d$ (p. 71)
- $\text{prox}_{\rho f}(x)$ _____ Proximal operator of function ρf at x (p. 58)
- \mathbb{Q}^d _____ Set of rationals in \mathbb{R}^d (p. 42)
- \mathcal{Q} _____ Operator class (p. 18)
- $\mathcal{Q}_{-L, L}(\mathbb{R}^d)$ _____ Class of L -Lipschitz operators on \mathbb{R}^d (p. 79)
- $\mathcal{Q}_{\mu, L}(\mathbb{R}^d)$ _____ Class of μ -strongly monotone
 L -Lipschitz operators on \mathbb{R}^d (p. 91)
- \mathbb{R} _____ Real numbers (p. 1)
- $\bar{\mathbb{R}}$ _____ Real extended numbers (p. 11)
- $\mathbb{R}_k[x_1, \dots, x_n]$ _____ Set of n -dimensional real-valued (p. 80)
polynomials of degree k
- S^n _____ Cone of symmetric matrices in $\mathbb{R}^{n \times n}$ (p. 23)
- $\mathcal{S}_{M,+}$ _____ Class of univariate self-concordant functions (p. 122)
- S_+^n _____ Cone of positive semidefinite definite matrices in S^n (p. 23)
- $\mathcal{T}_{M,+}$ _____ Class of univariate quasi self-concordant functions (p. 122)
- $\Sigma_{N, \mathbb{R}}$ _____ Covariance class (p. 193)
- $\Sigma_{N, 2^p}$ _____ Super-symmetric covariance class (p. 199)
- $\Sigma_{N, 2^N}^*$ _____ Completely anti-correlated covariance matrices (p. 203)
- $\text{Tr}(\cdot)$ _____ Matrix's trace (p. 22)
- $\omega(\cdot)$ _____ Function satisfying $\lim_{t \rightarrow 0} \frac{\omega(t)}{t} = 0$ (p. 33)
- x_* _____ Minimizer of function f (p. 1)
- $(x_k)_{k=0,1,\dots,N}$ _____ Sequence of iterates of a method (p. 5)

List of acronyms

| | |
|---------------|---------------------------------------|
| CNM | Cubic regularized Newton Method |
| DNM | Damped Newton Method |
| EG | ExtraGradient method |
| FSLFOM | Fixed-Step Linear First-Order Methods |
| GD | Gradient Descent |
| GNM | Gradient regularized Newton Method |
| IQC | Integral Quadratic Constraints |
| NM | Newton's Method |
| OG | Optimistic Gradient descent |
| PEP | Performance Estimation Problem |
| RCD | Randomized Coordinate descent |
| SDP | SemiDefinite Program |
| SGD | Stochastic Gradient Descent |
| SOS | Sum Of Squares |

1

Introduction

Optimization, which consists in finding the decision that minimizes or maximizes a given objective, is omnipresent in both nature and human activity. Many biological and physical processes can be interpreted through the lens of optimization—for instance, bird formations that reduce total energy expenditure, or physical systems that settle into configurations minimizing potential energy. On a societal level, optimization is, for better or worse, deeply embedded in our daily lives, from minimizing costs and maximizing efficiency to scheduling.

Formally, an optimization problem can be expressed as

$$\min_{x \in C} f(x), \tag{1.1}$$

where x is the *decision variable* encoding the possible actions or configurations, $f : C \rightarrow \mathbb{R} \cup \{+\infty\}$ is the *objective function* to be minimized, and C is the *feasible set*, representing the constraints that limit the decisions. The goal is to find a global *minimizer* $x_* \in C$ such that $f(x_*) \leq f(x)$ for all $x \in C$, when such a minimizer exists.

Despite its apparent simplicity, solving (1.1) analytically is, in most practical cases, difficult. This has led to the development of *optimization methods*, or algorithms, which iteratively approximate solutions to (1.1). Such algorithms are typically effective only for certain classes of problems, exhibiting specific structure. Indeed, designing a method that guarantees

to efficiently solve (1.1) without further assumptions on its structure is out of reach [Nes18, Chapter 1].

Defining and understanding what assumptions make a problem class desirable, designing algorithms suited to such problem classes, and analyzing their *performance*, i.e., whether the method always converges to a solution of (1.1), and if so, how fast, are three central challenges in optimization theory. This thesis considers some of these challenges from two perspectives: on one hand, it seeks to better characterize the structural properties defining favorable problem classes; on the other, it investigates the convergence behavior of iterative methods on these problem classes.

Before outlining the main contributions of this research, we provide a high-level overview of a few problem classes considered, some methods studied, and the tools used to analyze their performance.

Problem classes. For (1.1) to be solvable, one must add assumptions on both f and C . Such assumptions balance between (i) the theoretical guarantees they provide, i.e., it should be possible to design optimization methods that perform efficiently given these assumptions; and (ii) the scope of the assumptions, i.e., they should be satisfied by as many applications as possible.

For instance, classical function properties include (strong) *convexity* (Definition 2.2, see, e.g., [Roc70, Boy04]) and *L-smoothness*, see (2.5). Convexity enables global optimality analysis through local conditions. For example, when $C = \mathbb{R}^d$, any *stationary* point of a convex function, i.e., any point x satisfying $\nabla f(x) = 0$, a necessary condition for optimality, is also a global minimizer [Nes18, Theorem 2.1.2]. Strong convexity further guarantees uniqueness of the minimizer, at the cost of being more restrictive in the applications it encompasses [Nes18, Section 2.1.3]. Smoothness bounds how rapidly the gradient of a function can vary. Additionally, one may also require regularity of higher-order information, e.g., by ensuring Lipschitz continuity of the Hessian $\nabla^2 f(x)$, see, e.g., [Nes18, Chapter 4].

These properties can be combined to define function classes with strong theoretical guarantees. In particular, this thesis often considers the class of L -smooth (strongly) convex functions over \mathbb{R}^d , denoted $\mathcal{F}_{\mu,L}$, even though broader classes will be addressed in Chapters 4, 5, and 7.

First-order black-box methods. To solve minimization problems as (1.1), one classically uses black-box iterative methods, which, at each iteration, rely on a pointwise evaluation of the objective function, obtained through

an *oracle*, to update their approximate solution of (1.1). The outputs of the oracle form the only information on f available to the method, which thus cannot directly leverage the function's structure or its closed-form expression, see, e.g., [Nes18, Section 1.1.2]. This allows analyzing methods in a generic way, valid for all instances of a problem class, rather than on a case-by-case basis.

Such methods can be classified by the order of information returned by the oracle. For instance, first-order black-box methods have access, at each iteration, to function and gradient evaluations, $f(x)$ and $\nabla f(x)$. The simplest such method is the *gradient method*, first introduced in [C⁺47], and widely analyzed afterwards, see, e.g., [Pol63], whose update at iteration k is given by

$$x_{k+1} = x_k - \alpha_k \nabla f(x_k), \quad (\text{GD})$$

for some stepsize $\alpha_k > 0$. On favorable problem classes, and for wisely chosen stepsizes, (GD) converges to a stationary point of the objective function, which, depending on the problem class, may then automatically be a global minimizer of (1.1).

First-order methods are widely used due to their scalability, that is, their iterations are computationally cheap, especially in large-scale problems where higher-order information (e.g., Hessians) is either unavailable or computationally prohibitive to obtain and deal with. The price for this scalability is slower convergence, i.e., first-order methods typically require more iterations to reach the same level of accuracy than second-order methods. This slower convergence is fundamental, and holds even for the fastest first-order methods, see, e.g., [Nes83]. For further information on first-order methods, see, e.g., [Bec17].

Second-order black-box methods. Second-order optimization methods use, at each iteration, not only gradient information $\nabla f(x)$, but also curvature information provided by the Hessian $\nabla^2 f(x)$, to update their approximate solution to (1.1). The most classical example is *Newton's method*, first introduced by Newton and Raphson in the 18th century, see, e.g., [Pol07], which updates the iterate according to

$$x_{k+1} = x_k - [\nabla^2 f(x_k)]^{-1} \nabla f(x_k), \quad (\text{NM})$$

when the Hessian is invertible. Second-order methods typically enjoy a faster local convergence than first-order methods. In particular, on sufficiently regular objective functions, their convergence is *quadratic* near the minimizer of the function. On the downside, the computation and inversion of the Hessian may be prohibitively expensive in high dimensions, making these methods unpractical on large-scale applications; and simple methods such as Newton’s method may fail on simple problems [Nes18, Chapter 4]. In recent years, however, many variants to Newton’s method have been designed to overcome this issue, see, e.g., [NP06, Pol09, HKP⁺22, Mis23, DN24].

Stochastic methods. In some settings, to solve (1.1), it is convenient to consider *stochastic* optimization methods, which, to iteratively minimize the objective function f , rely on *noisy* evaluations of f and its derivatives. The simplest such method, counterpart to the deterministic (GD), is the stochastic gradient method (SGD), introduced in [RM51], whose iteration is given by

$$x_{k+1} = x_k - \alpha_k(\nabla f(x_k) + \varepsilon_k), \quad (\text{SGD})$$

where ε_k is a random variable, depending on the setting. Stochastic methods are used when (i) the exact gradient of f is unavailable, e.g., when f is defined as an expectation,

$$f(x) := \mathbb{E}_{\xi}[F_{\xi}(x)],$$

as in supervised learning [BCN18], or (ii) computing the exact gradient is costly, as in finite-sum problems

$$f(x) = \frac{1}{n} \sum_{i=1}^n f_i(x),$$

where n is large [BB07]. In such cases, one typically estimates $\nabla f(x)$ using a stochastic gradient $\nabla f_i(x)$ based on a randomly sampled index i , which is approximately n times cheaper to compute than $\nabla f(x)$. Given their scalability, stochastic first-order methods, such as SGD and its many variants, are central to modern machine learning and have been extensively studied, see, e.g., [Bor08, BCN18, GG23, GLQ⁺19, Net19].

Worst-case performance guarantees. To assess the efficiency of an optimization method, a common approach is to analyze its *worst-case* performance over a given problem class. That is, one seeks upper bounds on a performance criterion, typically the objective value accuracy after a fixed number of iterations N , e.g., $f(x_N) - f(x_*)$, where x_* is a minimizer or stationary point of f , that hold for *any* instance in the problem class. While such guarantees may be pessimistic in practice, as not all instances in the class induce the worst behavior of the method, they provide strong theoretical insights into the behavior of an algorithm and its bottlenecks, i.e., on which problem instances the method performs slowly. Moreover, worst-case guarantees allow estimating, a priori, the number of iterations required to achieve a given accuracy, independently of the problem instance. For these reasons, we adopt this framework in the present thesis. Other performance assessments include, e.g., averaged-case [Bor87], or smoothed [ST04] performance guarantees.

Formally, given a method \mathcal{M} , a function class \mathcal{F} , an initial condition C_{init} controlling the initial distance to optimality, and a performance criterion \mathcal{P} , depending on the objective function and the iterates $(x_k)_{k=0,1,\dots,N}$ of \mathcal{M} , we aim to obtain upper bounds on

$$\begin{aligned} \max_{f \in \mathcal{F}, x_0 \in \mathbb{R}^d, x_* \in \mathbb{R}^d} \mathcal{P}(f, (x_k)_{k=0,1,\dots,N}) \text{ s.t. } (x_k)_{k=0,1,\dots,N} \text{ is generated by } \mathcal{M}, \\ C_{\text{init}}(f, x_0) \text{ is satisfied, and} \\ x_* \text{ is a stationary point of } f. \end{aligned} \tag{1.2}$$

On stochastic methods, worst-case guarantees may additionally depend on the distribution of the noise, i.e., they should hold for any distribution satisfying given assumptions, and the objective function in (1.2) becomes the expected value of \mathcal{P} with respect to the randomness in the method.

We seek guarantees that are as accurate as possible, and call a bound *tight* when it is attained by some instance in the problem class, that is, the solution f to (1.2). Tight worst-case guarantees enable fair comparisons between methods and efficient tuning of methods, by selecting the method's parameters yielding the best theoretical performance. By contrast, tuning or comparing methods on the basis of upper bounds on the actual performance may lead to sub-optimal or misguided choices, as illustrated on Figure 1.1.

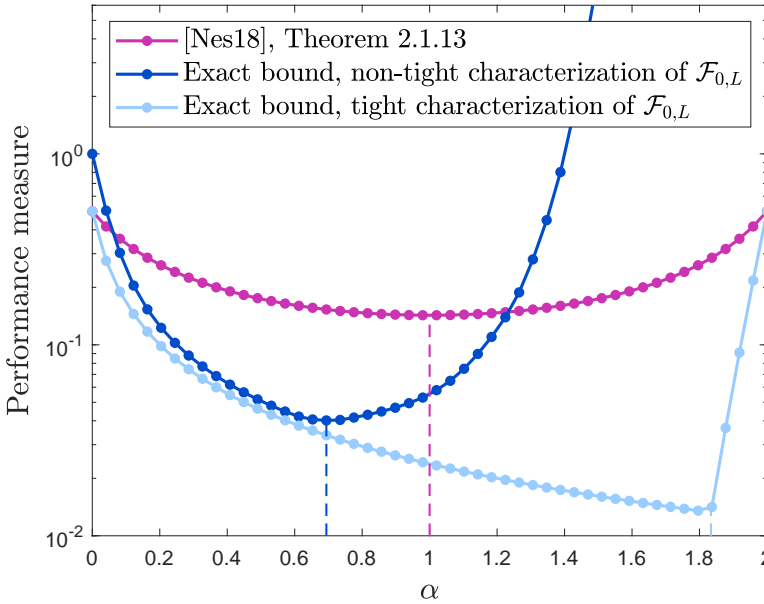


Fig. 1.1 Comparison of performance analyses of 10 iterations of (GD) with constant stepsize α on $\mathcal{F}_{0,L}$, with $L = 1$. The initial condition is $\|x_0 - x_*\| \leq 1$. Displayed is the evolution of the performance measure $f(x_{10}) - f(x_*)$, as a function of the stepsize α , as predicted by (i) the theoretical bound from [Nes18, Theorem 2.1.13] (pink), (ii) the best bound possible (PEP-based), relying on the non-tight representation of $\mathcal{F}_{0,L}$ (2.7) (dark blue), and (iii) the exact (PEP-based) bound, relying on the exact characterization of $\mathcal{F}_{0,L}$ (2.10) (light blue).

Computer-aided worst-case performance analysis Obtaining worst-case guarantees, and ideally, tight ones, is generally challenging when done manually. Fortunately, this process can be automated via computer-aided techniques as, e.g., the *Integral Quadratic Constraints* (IQC) [LRP16], or the *Performance Estimation Problem* (PEP) [DT14, THG17c] frameworks. The latter, for instance, computes exact worst-case performance guarantees by solving the optimization problem (1.2), thereby identifying worst-case instances within a given problem class. Its core idea is to reformulate this infinite-dimensional problem as a finite-dimensional one: a maximization over a set of quantities S , encoding the method's iterates, a minimizer of the objective f , and evaluations of f , its gradient, and possibly its Hes-

sian at those points, under the constraint that S is consistent with an actual function in \mathcal{F} . This requires *interpolation conditions* for \mathcal{F} , i.e., necessary and sufficient conditions that a dataset S must satisfy to be consistent with a function in \mathcal{F} . These interpolation conditions serve as *discrete* characterizations of the properties defining \mathcal{F} , such as convexity or smoothness, but verified only on a finite set of points.

Whenever interpolation conditions are available and tractable, solving the corresponding PEP formulation, which can be done efficiently for a wide range of first-order problems, yields tight performance bounds. By contrast, when using only *necessary* conditions to ensure that the set S is consistent with a function in \mathcal{F} , the resulting bounds may hold for a larger problem class than originally intended. In this case, the guarantees are a priori conservative, as the worst-case instance identified might not belong to the initial problem class. Interpolation conditions, tracing back to [Kir34, Tie15, Whi34] will be discussed in Chapter 2. The PEP framework and other computer-aided approaches to performance analysis, see, e.g., [LRP16], will be presented in Chapter 3.

From a broader perspective, proofs of worst-case guarantees, whether manual or PEP-based, generally amount to combining algebraic relations that encode the properties of the problem class and the dynamics of the method, in order to bound a performance criterion. Conservatism in the resulting bound may arise from two main sources: (i) *weak* algebraic representations that fail to fully capture the function class properties, and (ii) a suboptimal combination of these relations. Interpolation conditions address the first issue by providing exact, discrete characterizations of the function class. The PEP framework addresses the second by optimally combining the inequalities it is given to produce the tightest possible bound under the assumptions.

To emphasize the importance of not only optimally combining inequalities but also relying on accurate characterizations of function classes, Figure 1.1 compares several performance guarantees for (GD) with constant stepsizes $\alpha \in [0, \frac{2}{L}]$, applied to the class $\mathcal{F}_{0,L}$ of L -smooth convex functions. Specifically, it displays: (i) an analytical upper bound from [Nes18, Theorem 2.1.13], (ii) a PEP-based bound using an inexact description of $\mathcal{F}_{0,L}$, and (iii) the exact convergence guarantee. The optimal stepsizes suggested by these bounds differ significantly, and notably, the one inferred from the inexact PEP-based analysis leads to the worst actual performance.

1.1 Outline and contributions

This thesis aims, from one side, at better understanding and obtaining interpolation conditions for function and operator classes; and, from another side, at extending the PEP framework to the analysis of several optimization settings. In particular, combining these two aspects, we *tightly* analyze (i) broader first-order problem classes, including non-convex settings; (ii) second-order methods in the univariate case; and (iii) stochastic first-order algorithms. More precisely, this work is structured as follows.

Preliminaries

Chapter 2 introduces function and operator interpolation, and reviews the state of the art and open challenges in this field.

Chapter 3 introduces the core elements of the PEP framework and other computer-aided techniques for performance analysis, along with their current developments.

Part I: An algebraic approach to function and operator interpolation

Chapter 4 introduces a constraint-based approach to interpolation, setting aside analytic properties of function or operator classes to work at an algebraic level. This approach is then used to derive interpolation conditions for weakly convex functions with bounded subgradients, which allows, numerically, to tightly analyze first-order methods such as the subgradient method on this class. *This chapter is based on [RH24].*

Chapter 5 builds on Chapter 4 to propose a refinement procedure that iteratively strengthens any given description of a function or operator class into, eventually, an interpolation condition. The procedure is applied to strengthen existing descriptions of several function and operator classes, including smooth functions satisfying a Łojasiewicz condition, convex blockwise smooth functions, Lipschitz (strongly) monotone operators, and uniformly convex functions. The benefits of relying on the improved descriptions are illustrated on numerical examples. *This chapter is based on [RHT25]. It results from a collaboration with Adrien Taylor.*

Part II: Interpolation conditions for second-order univariate optimization.

Chapter 6 presents a generic method to obtain interpolation conditions for a broad range of *univariate* function classes. *This chapter is based on [RBHG25]. It results from a collaboration with Nizar Bousselmi.*

Chapter 7 applies the method introduced in Chapter 6 to derive interpolation conditions for a general class of univariate functions characterized by second-order properties, which includes self-concordant functions and (strongly convex) functions with Lipschitz-continuous Hessians. These conditions are then used to analytically and numerically study the exact worst-case performance of various second-order methods in the univariate setting. *This chapter is based on [RBHG25]. It results from a collaboration with Nizar Bousselmi.*

Part III: Interpolation conditions for stochastic first-order optimization.

Chapter 8 develops a methodology, embedded within the PEP framework, for deriving worst-case performance guarantees on stochastic first-order methods. Specifically, given a first-order method, a function class, and a noise model with prescribed expectation and variance properties, the framework proposes a hierarchy of optimization problems of increasing size, yielding increasingly strong convergence guarantees. The largest formulation, whose size grows exponentially with the number of iterations, provides tight guarantees. This framework is then used to analyze, both numerically and analytically, the stochastic gradient descent under various noise models, including unstructured noise with bounded variance, finite-sum problems, and block-coordinate methods. *This chapter is based on [RCH25].*

Finally, **Chapter 9** summarizes the research outcomes and discusses the perspectives of this thesis.

List of publications

1.1 Journal papers

- [RH24] Anne Rubbens and Julien M. Hendrickx. **"A constraint-based approach to function interpolation, with application to performance estimation for weakly convex optimization."**, *arXiv preprint*, 2024.
- [RHT25] Anne Rubbens, Julien M. Hendrickx, and Adrien Taylor. **"A constructive approach to strengthen algebraic descriptions of function and operator classes"**. *arXiv preprint*, 2025.
- [RBHG25] Anne Rubbens, Nizar Bousselmi, Julien M. Hendrickx, and François Glineur. **"Performance estimation of second-order optimization methods on classes of univariate functions"**. *arXiv preprint*, 2025.
- [RCH25] Anne Rubbens, Sébastien Colla, and Julien M. Hendrickx. **"Computer-aided analyses of stochastic first-order methods, via interpolation conditions for stochastic optimization."** *arXiv preprint*, 2025.

1.1 Abstracts and conference proceedings

- [RBCH23] Anne Rubbens, Nizar Bousselmi, Sébastien Colla and Julien M. Hendrickx, **"Interpolation Constraints for Computing Worst-Case Bounds in Performance Estimation Problems"**, in Proceedings of the 62th IEEE Conference on Decision and Control, 2023.

2

Interpolation conditions for function and operator classes

This chapter introduces the notion of interpolation conditions on the example of L -smooth convex functions, and provides a survey of related literature.

Notation Throughout, \mathbb{R} denotes the real numbers, $\bar{\mathbb{R}} := \mathbb{R} \cup \{+\infty\}$ denotes the real extended numbers, see, e.g., [Bec17, Chapter 2.1], and \mathbb{R}^d denotes the vector space of real d -tuples. We let $|\cdot|$ denote the absolute value, and consider the standard dot product $\langle \cdot, \cdot \rangle$, where given $x, y \in \mathbb{R}^d$, $\langle x, y \rangle = x^\top y = \sum_{i=1}^d x_i y_i$, and induced norm $\|\cdot\| = \sqrt{\langle \cdot, \cdot \rangle}$. Given $N \in \mathbb{N}$, we let $[N] = \{1, \dots, N\}$. Given a function $f : C \rightarrow \bar{\mathbb{R}}$, we denote by $\text{dom} f := \{x \in C : f(x) < +\infty\}$, the effective domain of f , and consider proper functions, i.e., $\text{dom} f \neq \emptyset$.

2.1 Global characterizations of function classes

Function classes are typically defined via algebraic relations that hold throughout the function's domain. These relations may not be unique, i.e., several sets of algebraic relations may equivalently characterize a single function class. We show some examples of these characterizations on convex and L -smooth functions.

Convex functions We first recall the definition of a convex set.

Definition 2.1 (Convex set). A set $C \subset \mathbb{R}^d$ is said to be *convex* if, $\forall x, y \in C$ and $\alpha \in [0, 1]$, $\alpha x + (1 - \alpha)y \in C$.

Convex functions, widely investigated in optimization theory, see, e.g., [Boy04, Roc70, Nes18], can be defined in several ways. The maybe simplest one uses the same quantifier as in the definition of convex sets.

Definition 2.2 (Convex functions). Let $C \subseteq \mathbb{R}^d$. A function $f : C \rightarrow \bar{\mathbb{R}}$ is *convex* if $\text{dom } f$ is convex, and $\forall x, y \in \text{dom } f$, and $\forall \alpha \in (0, 1)$,

$$f(\alpha x + (1 - \alpha)y) \leq \alpha f(x) + (1 - \alpha)f(y). \quad (2.1)$$

Alternatively to Definition 2.2, convex functions can be defined via their *subdifferential* $\partial f(x)$. The subdifferential generalizes the notion of a gradient to non-differentiable functions, and is defined as the set of all *subgradients* of f at x , when $x \in \text{dom } f$, that is, vectors $v_x \in \mathbb{R}^d$ satisfying

$$v_x \in \partial f(x) \Leftrightarrow f(y) \geq f(x) + \langle v_x, y - x \rangle, \quad \forall y \in C. \quad (2.2)$$

At points $x \notin \text{dom } f$, we set $\partial f(x) = \emptyset$. When f is differentiable at x , the subdifferential reduces to the gradient, i.e., $\partial f(x) = \{\nabla f(x)\}$; see, e.g., [RWW09, Exercise 8.8].

Definition 2.3 (Convex functions: subdifferential definition, from [Tay17], Theorem 2.14). Let $C \subseteq \mathbb{R}^d$. A function $f : C \rightarrow \bar{\mathbb{R}}$ is *convex* if and only if $\text{dom } f$ is convex, and $\forall x \in \text{relint}(\text{dom } f)$, $\partial f(x) \neq \emptyset$.

By Definition 2.3, a differentiable function $f : C \rightarrow \mathbb{R}$, i.e., $\nabla f(x)$ exists at all $x \in C$, is convex if and only if, $\forall x, y \in C$,

$$f(x) \geq f(y) + \langle \nabla f(y), x - y \rangle. \quad (2.3)$$

Variations to convex functions include μ -strongly convex functions, that is functions $f : C \rightarrow \bar{\mathbb{R}}$ such that, for some $\mu > 0$, $f(\cdot) - \frac{\mu}{2} \|\cdot\|^2$ is convex.

L-smooth functions Aside to convexity, widely used properties of functions involve regularity of the function and its gradient.

Definition 2.4 (Lipschitz continuous function). Let $C \subseteq \mathbb{R}^d$, and $L \geq 0$. A function $f : C \rightarrow \mathbb{R}$ is *L-Lipschitz continuous* if, for all $x, y \in C$,

$$|f(x) - f(y)| \leq L\|x - y\|. \quad (2.4)$$

Lipschitz continuity can be lifted to derivatives of a function. In particular, a function f is *L-smooth* for some $L \geq 0$ if it is differentiable on its domain, and its gradient is *L-Lipschitz continuous*, i.e., for all $x, y \in C$,

$$\|\nabla f(x) - \nabla f(y)\| \leq L\|x - y\|. \quad (2.5)$$

Alternatively, when $C = \mathbb{R}^d$, *L-smoothness* can be characterized via a relation involving function values, obtained by integration of the Lipschitz continuity of the gradients, see, e.g., [Nes18, Lemma 1.2.3].

Definition 2.5 (*L-smooth functions*). Given $L \geq 0$, we say $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is *L-smooth* if f is everywhere differentiable and, $\forall x, y \in \mathbb{R}^d$,

$$|f(x) - f(y) - \langle \nabla f(y), x - y \rangle| \leq \frac{L}{2} \|x - y\|^2. \quad (2.6)$$

Function class $\mathcal{F}_{\mu,L}(C)$ Properties of convexity and smoothness can also be combined. Given $C \subseteq \mathbb{R}^d$, and $0 \leq \mu \leq L \leq +\infty$, we denote by $\mathcal{F}_{\mu,L}(C)$ the class of proper μ -strongly convex *L-smooth* functions with domain C , where $L = +\infty$ refers to non-smooth functions. Again, setting $C = \mathbb{R}^d$, the class $\mathcal{F}_{\mu,L}(\mathbb{R}^d)$ ($L \neq +\infty$), referred to as $\mathcal{F}_{\mu,L}$ when clear from the context, can be defined in numerous equivalent ways, see, e.g., [Nes18, Section 2.1.1]. Arguably the most natural one consists in juxtaposing a (strong) convexity condition (2.3), with Lipschitz continuity of the gradients, i.e., (2.5). That is, $f : \mathbb{R}^d \rightarrow \mathbb{R}^d$ belongs to $\mathcal{F}_{\mu,L}$ if and only if it is everywhere differentiable and it satisfies, $\forall x, y \in \mathbb{R}^d$,

$$\begin{cases} f(x) \geq f(y) + \langle \nabla f(y), x - y \rangle + \frac{\mu}{2} \|x - y\|^2 \\ \|\nabla f(x) - \nabla f(y)\| \leq L\|x - y\|. \end{cases} \quad (2.7)$$

Alternatively, one can rely on (2.6) instead of (2.5) to characterize smoothness. That is, $f : \mathbb{R}^d \rightarrow \mathbb{R}^d$ belongs to $\mathcal{F}_{\mu,L}$ if and only if it is everywhere differentiable and, $\forall x, y \in \mathbb{R}^d$,

$$\begin{cases} f(x) \geq f(y) + \langle \nabla f(y), x - y \rangle + \frac{\mu}{2} \|x - y\|^2 \\ f(x) \leq f(y) + \langle \nabla f(y), x - y \rangle + \frac{L}{2} \|x - y\|^2. \end{cases} \quad (2.8)$$

One can further show that $\mathcal{F}_{\mu,L}$ ($\mu \neq L$) is equivalently defined by the following single condition, when imposed on all $x, y \in \mathbb{R}^d$, see, e.g., [Nes18, Theorem 2.1.12]:

$$\begin{aligned} f(x) \geq f(y) + \langle \nabla f(y), x - y \rangle & \quad (2.9) \\ + \frac{L}{2(L - \mu)} \left(\frac{\|\nabla f(x) - \nabla f(y)\|^2}{L} + \mu \|x - y\|^2 \right. \\ & \left. - \frac{2\mu}{L} \langle \nabla f(x) - \nabla f(y), x - y \rangle \right). \end{aligned}$$

When $\mu = 0$, (2.9) reduces to

$$f(x) \geq f(y) + \langle \nabla f(y), x - y \rangle + \frac{1}{2L} \|\nabla f(x) - \nabla f(y)\|^2. \quad (2.10)$$

Finally, twice differentiable functions belong to $\mathcal{F}_{\mu,L}$ ($\mu \neq L$) if only if they satisfy, $\forall x \in \mathbb{R}^d$ [Nes18, Theorem 2.1.6],

$$\mu \leq \nabla^2 f(x) \leq L, \quad \forall x, y \in \mathbb{R}^d. \quad (2.11)$$

2.2 Discrete characterizations of function classes

All characterizations of function classes introduced in Section 2.1 consist of pairwise algebraic relations, required to hold at any pair in the effective domain of functions in the class. We call such characterizations *global*. On the contrary, we seek for *discrete* characterizations of these function classes \mathcal{F} , i.e., sets of conditions answering, for finite sets $\mathcal{X} \subset \mathbb{R}^d$:

What conditions must a function f satisfy on \mathcal{X} to ensure existence of an extension of f that is identically equal to f on \mathcal{X} and belongs to \mathcal{F} ?

The need for discrete characterizations arises in any situation where one

manipulates or has access to a finite dataset $S = \{(x_i, f_i, g_i)\}_{i \in [N]}$, and requires it to be \mathcal{F} -interpolable, that is, consistent with some function $f \in \mathcal{F}$, see Figure 2.1 for an illustration.

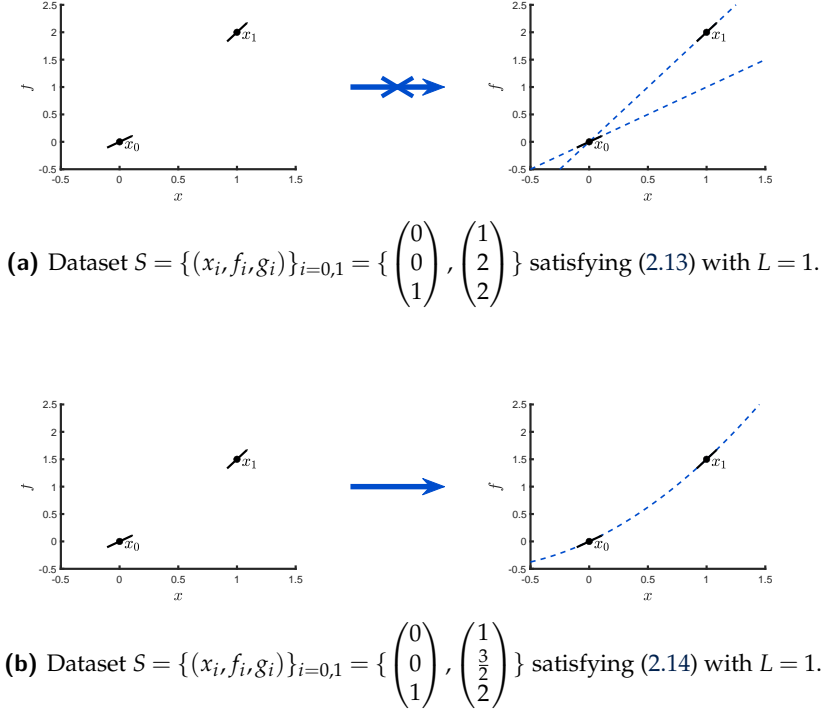


Fig. 2.1 From [Tay17, Figure 3.1]. Discrete interpolation, or extension, problem. Given $S = \{(x_i, f_i, g_i)\}_{i \in [N]}$, what conditions must the dataset satisfy to ensure consistency with some function in \mathcal{F} ? In (a), a dataset satisfying (2.13) is not consistent with any actual function in $\mathcal{F}_{0,L}$, where $L = 1$. Indeed, by convexity, any function in $\mathcal{F}_{0,L}$ interpolating S should remain above its linear under-approximations, displayed in dotted lines. This creates a non-differentiability at x_0 , incompatible with L -smoothness of the interpolating function. In (b), a dataset satisfying the interpolation condition (2.14) is $\mathcal{F}_{0,L}$ -interpolable, where $L = 1$, since consistent with the L -smooth convex function $f(x) = \frac{x^2}{2} + x$.

Definition 2.6 (\mathcal{F} -interpolability). Let $N \in \mathbb{N}$, \mathcal{F} a function class, and a set $S = \{(x_i, f_i, g_i)\}_{i \in [N]} \in (\mathbb{R}^d \times \mathbb{R} \times \mathbb{R}^d)^N$. We say S is \mathcal{F} -interpolable if and

only if

$$\exists f \in \mathcal{F} : f(x_i) = f_i, \text{ and } g_i \in \partial f(x_i), \quad \forall i \in [N]. \quad (2.12)$$

We refer by *interpolation conditions* for \mathcal{F} , to sets of conditions whose satisfaction by *any* dataset S is necessary and sufficient for \mathcal{F} -interpolability of S . Such conditions provide an exact discrete characterization of \mathcal{F} . A natural way to obtain discrete characterizations is to *discretize* global characterizations, that is, to impose them on a smaller subset of points, since any of these discretizations is *necessarily* satisfied by \mathcal{F} -interpolable datasets. For instance, consider $\mathcal{F}_{0,L}$, defined equivalently, on the whole domain, by (2.7) with $\mu = 0$ and (2.10). Any dataset $S = \{(x_i, f_i, g_i)\}_{i \in [N]}$ necessarily satisfies the discretized version of these characterizations, e.g., $\forall i, j \in [N]$,

$$\begin{cases} f_i & \geq f_j + \langle \nabla g_j, x_i - x_j \rangle \\ \|g_i - g_j\| & \leq L \|x_i - x_j\|, \end{cases} \quad (2.13)$$

and

$$f_j \geq f_i + \langle g_i, x_i - x_j \rangle + \frac{1}{2L} \|g_i - g_j\|^2. \quad (2.14)$$

However, the equivalence between the constraints no longer holds when they are imposed on a finite number of pairs only. In particular, (2.14) is stronger than (2.13), since one can find datasets $S = \{(x_i, g_i, f_i)\}$ which satisfy (2.13), but are not consistent with an actual function in $\mathcal{F}_{0,L}$, see, e.g., Figures 2.1 and 2.2. While it is a necessary and sufficient global characterization of $\mathcal{F}_{0,L}$, Condition (2.7) is thus not sufficient for interpolation. On the contrary, it has been shown [THG17c, Theorem 4] that (2.10) (and, more generally, (2.9)) is a necessary and sufficient interpolation condition for $\mathcal{F}_{0,L}$ (resp. $\mathcal{F}_{\mu,L}$).

Proposition 2.1 ([THG17c], Theorem 4: interpolation conditions for $\mathcal{F}_{\mu,L}$). *Given $0 \leq \mu < L \leq +\infty$ and $N \in \mathbb{N}$, let $\mathcal{F}_{\mu,L}$ and a set $S = \{(x_i, f_i, g_i)\}_{i \in [N]} \subset (\mathbb{R}^d \times \mathbb{R} \times \mathbb{R}^d)^N$. Then, S is $\mathcal{F}_{\mu,L}$ -interpolable if and only if, $\forall i, j \in [N]$,*

$$\begin{aligned} f_j & \geq f_i + \langle g_i, x_j - x_i \rangle \\ & + \frac{L}{2(L - \mu)} \left(\frac{\|g_i - g_j\|^2}{L} + \mu \|x_i - x_j\|^2 - \frac{2\mu}{L} \langle g_i - g_j, x_i - x_j \rangle \right). \end{aligned} \quad (2.15)$$

For pedagogical purposes, we introduced discrete characterizations on first-order function classes, i.e., classes whose descriptions involve function and gradient evaluations. But, all concepts can be straightforwardly

extended to broader settings, as will be illustrated in Chapter 6.




Fig. 2.2 Given $L = 1$, $(x_*, f_*, g_*) = (0, 0, 0)$, $(x_0, f_0, g_0) = (1, \frac{1}{4}, \frac{1}{2})$ and $x_1 = \frac{1}{2}$, allowed region for (f_1, g_1) according to (2.13) and (2.14). The region allowed by (2.14) is included in and significantly smaller than the one allowed by (2.13). In particular, no boundary is shared by the two regions. This shows that, despite the equivalence between these constraints when imposed on all pairs (x, y) , Condition (2.13) is weaker than (2.14). In particular, the points in the difference between the two regions do not correspond to any actual function $f \in \mathcal{F}_{0,L}$. In addition, when solving optimization problems such as $\max_{f \in \mathcal{F}_{0,L}} f(x_1) - f_*$ s.t. $x_1 = x_0 - \frac{1}{L} \nabla f(x_0)$, using one or the other characterization to translate $f \in \mathcal{F}_{0,L}$ into conditions on $S = \{(x_i, f_i, g_i)\}_{i=0,1,*}$ heavily modifies the resulting bound.

2.3 Operator classes and their discrete characterizations

Some optimization methods, rather than minimizing functions, seek to solve *variational inequalities* [Fic64, LS67, Sta64], i.e., to solve

$$\text{Find } x_* \in C : 0 \in Q(x_*), \tag{VI}$$

where $Q : \mathbb{R}^d \rightrightarrows \mathbb{R}^d$ is a (possibly set-valued) operator. The notation \rightrightarrows indicates that Q may assign multiple values in \mathbb{R}^d to each $x \in \mathbb{R}^d$. We refer to any of these values as $q_x \in Q(x)$.

Operators of interest include, e.g., *linear operators*, i.e., $Q(x) = Mx$ for some $M \in \mathbb{R}^{d \times d}$, and *subdifferentials*, i.e., $Q(x) = \partial f(x)$, for some finite-valued convex function f . In this second case, looking for a minimizer of f is equivalent to solving (VI).

As for function minimization, methods to solve (VI) depend on properties satisfied by operators Q on their effective domain, a few of which we state here, see, e.g., [RB16] for further reference. The pendant to convexity is denoted by (*strong*) *monotonicity*, i.e., an operator $Q : \mathbb{R}^d \rightrightarrows \mathbb{R}^d$ is μ -*strongly monotone* if, $\forall x, y \in \mathbb{R}^d$,

$$\langle q_x - q_y, x - y \rangle \geq \mu \|x - y\|^2, \quad \forall q_x \in Q(x), q_y \in Q(y). \quad (2.16)$$

When $\mu = 0$, such operators are referred to as *monotone* operators. A monotone operator Q is said to be *maximal* if its graph is not properly contained in any other monotone operator, i.e., if there exists no other monotone operator \bar{Q} such that

$$\{(x, y) \in \mathbb{R}^d \times \mathbb{R}^d \mid y \in Q(x)\} \subsetneq \{(x, y) \in \mathbb{R}^d \times \mathbb{R}^d \mid y \in \bar{Q}(x)\}.$$

Single-valued operators can be Lipschitz continuous, as, e.g., gradients of Lipschitz continuous functions, i.e., they satisfy, $\forall x, y \in \mathbb{R}^d$,

$$\|Q(x) - Q(y)\| \leq L \|x - y\|. \quad (2.17)$$

Finally, another important property of operators is their co-coercivity, i.e., satisfaction, $\forall x, y \in \mathbb{R}^d$, of

$$\langle q_x - q_y, x - y \rangle \geq \beta \|q_x - q_y\|^2, \quad \forall q_x \in Q(x), q_y \in Q(y). \quad (2.18)$$

Co-coercivity implies both monotonicity and $\frac{1}{\beta}$ -Lipschitz continuity, but the converse does not hold in general.

Discrete characterizations of operators. Similarly to function classes, we seek for *discrete* characterizations of these operator classes \mathcal{Q} , that is for interpolation conditions ensuring a dataset $S = \{(x_i, t_i)\}_{i \in [N]}$ is \mathcal{Q} -*interpolable*, where we slightly modify the definition of interpolability.

Definition 2.7 (\mathcal{Q} -interpolability). Let $N \in \mathbb{N}$, \mathcal{Q} an operator class, and a set $S = \{(x_i, q_i)\}_{i \in [N]} \in (\mathbb{R}^d \times \mathbb{R}^d)^N$. We say S is \mathcal{Q} -interpolable if and only if

$$\exists Q \in \mathcal{Q} : q_i \in Q(x_i), \quad \forall i \in [N]. \quad (2.19)$$

2.4 State of the art in interpolation conditions

Extension of properties. What we call in this work interpolation is also referred to as extension. The interest in extension conditions can be traced back to at least 1915, with [Tie15] considering the extension of a continuous function outside of its domain. The question was then generalized in Whitney’s seminal paper [Whi34], seeking for interpolation conditions to impose on a (possibly infinite) dataset, for it to be consistent (up to order k derivative) with a function defined everywhere, and continuous up to a certain order of derivatives. In a parallel way, [Kir34, Val43, McS34] investigated the extension of Lipschitz continuous (and Hölder continuous, i.e., $\|f(x) - f(y)\| \leq L\|x - y\|^\alpha$, $\alpha \in (0, 1)$) functions to the whole domain.

Building on these works, interpolation conditions were then obtained for a wide range of settings, depending on the properties to be satisfied everywhere, on the order of derivatives up to which the extension should interpolate the dataset, and on the structure of the initial domain. See, e.g., [Gla58], [Wel73, Theorem 2], and [Gru09, Theorem 2.6], for the extension of L -smooth functions; [BS01, Fef05], for the interpolation of *function values* by functions continuous up to given order of derivatives, [FIL17] for the extension of non-negative continuous functions; and [AM17a, AM17b, AM19, Yan14, DHLGL18] for the extension of convex (smooth) functions.

Most results in this literature differ from the approach to interpolability we are interested in. First, some focus on qualitative extensions where regularity constants may increase, see, e.g., [AM17b, Gla58], while we seek quantitative guarantees, essential in the context of establishing properties related to these constants. Second, some results require assumptions on the domain, such as its convexity [AM19, Yan14], whereas we only impose conditions on finite sets of iterates.

First-order interpolation conditions Interpolation conditions in the context of optimization were introduced in [THG17c] to tightly analyze first-order methods on first-order function classes, i.e., methods and classes

whose characterization involves function and (sub)gradients evaluations. In this context, initial functions and extensions are required to satisfy properties with the same parameters. In [THG17c] and building on this work, interpolation conditions were obtained for a wide range of function classes, including smooth (strongly) convex functions [THG17c, Theorem 4], also shown to hold for negative values of μ , the convexity parameter [RGP22, Theorem 3.1], smooth Lipschitz functions [THG17a, Theorem 3.5], strongly convex functions with bounded domain [THG17a, Theorem 3.4], indicator functions [THG17a, Theorem 3.6], generalized in [LG24, Theorems 3.1, 3.2, 3.3], difference-of-convex and relatively smooth convex functions [DTdB22, Corollaries 1,2], convex functions with quadratic upper bounds [GTD22, Theorem 2.6], and alternative conditions to smoothness and convexity [GEIGM22, Theorem 1]. For further reference, see, e.g., the partial list provided in the PEPit documentation [GMG⁺24].

Zero-th order interpolation conditions, and interpolation conditions for operator classes Operator interpolation can be viewed as function interpolation for (possibly multi-valued) zero-th order function classes, i.e., classes whose characterization involves function evaluations only. Interpolation conditions exist for (strongly) monotone [BC17, Theorem 20.21], co-coercive [RTBG20, Proposition 2], negative cocoercive [GTHG23, Theorem B.1], and Lipschitz operators [Kir34, Val43]. In all three cases, these interpolation conditions happen to be the discretized versions of their classical global definitions, i.e., (2.16), (2.18), and (2.17). However, no interpolation result is known for operators combining several of these properties (except for datasets of cardinality 2, [RTBG20, Proposition 4]). In addition, recently, interpolation conditions were derived for linear operators with additional properties [BHG24, Theorem 3.1,3.3].

Finally, interpolation constraints for first-order function classes \mathcal{F} that do not explicitly involve function values, in case one would not want to use these as variables, also exist, based on Rockafellar's cyclic monotonicity conditions [Roc70, Tay17]. They can be seen as operator interpolation for a subclass of operators, coherent as gradients of a function.

Interpolation conditions for higher-order properties To this day, not much has been achieved on the topic of interpolation conditions for function classes involving higher-order properties, e.g., functions whose second or third derivative is Lipschitz continuous. In particular, to the best of our

knowledge, no interpolation condition exists yet for any second-order function class. Some progress has been made in obtaining interpolation conditions for convex functions, when these are meant to interpolate a dataset up to their second derivative [Yan14], but these are sufficient conditions only, and require the dataset to cover a convex region.

2.5 Challenges in obtaining interpolation conditions

Obtaining interpolation conditions for given function/operator classes remains to this day a challenging problem. In particular, there exists no principled way to answer some fundamental questions.

Question 1. Given a function/operator class, how to guess candidate interpolation conditions for this class?

Question 2. How to assess a constraint's interpolability, i.e., how to prove it is either a non-sufficient condition, or an interpolation condition for a given function or operator class?

Presently, answering these questions requires case-by-case techniques, that often heavily rely on analytical properties of the function/operator class analyzed, or involve intricate geometrical constructions. These techniques include, to establish interpolability of a given set of conditions for a class \mathcal{F} , typically a discretization of one of its global characterizations, (i) explicitly exhibiting a function in \mathcal{F} interpolating any dataset satisfying these conditions, or (ii) proving there exists such function without constructing it, using, e.g., Zorn's Lemma [Gru09]. In addition, given a function class \mathcal{F} , another strategy answering both Questions 1 and 2 consists in establishing a relation between this class and another one, for which interpolation conditions are known, using analytical properties of the class of interest, thereby obtaining both a global and a discrete characterization of this class, see, e.g., [THG17b]. These case-by-case techniques have not yet answered all interpolation conditions of interest in optimization, e.g., many function and operator classes lack interpolation conditions, including, in particular, non-convex function classes and function classes with higher-order properties. In addition, it is not known how to obtain interpolation conditions for function classes whose domain is a subset of \mathbb{R}^d , whose need arises in constrained optimization, see, e.g. [Dro20].

3

Computer-aided computation of worst-case performance guarantees

Computer-aided techniques for automatically analyzing and designing first-order methods have in the last decade gained significant attention; see, e.g., [Tay24] for a recent survey. These include (i) the automatic computation of worst-case instances for a given method over a function class \mathcal{F} after N iterations, known as the Performance Estimation (PEP) framework [DT14, THG17c]; (ii) the automated construction of Lyapunov (or potential) functions to certify (sub)linear convergence by analyzing a small number of iterations [LRP16, TVSL18, TB19]; and (iii) constructive approaches for designing optimal or efficient methods tailored to specific function classes [DGVPR24, DT20, KF16, TD23]. This chapter introduces the key elements of these approaches and provides a review of the corresponding literature.

Notation We denote by $\mathbb{R}^{m \times n}$ the set of real-valued $m \times n$ matrices. Given $A, B \in \mathbb{R}^{m \times n}$, we let $\text{Tr}(A^\top B) = \sum_{i=1}^m \sum_{j=1}^n A_{ij} B_{ij}$. The cones of *symmetric* matrices, and of *positive semidefinite* matrices in $\mathbb{R}^{n \times n}$, are denoted

by, respectively

$$\begin{aligned}\mathcal{S}^n &:= \{A \in \mathbb{R}^{n \times n} : A = A^\top\}, \\ \mathcal{S}_+^n &:= \{A \in \mathcal{S}^n : A \succcurlyeq 0\},\end{aligned}$$

where $A \succcurlyeq 0 \Leftrightarrow x^\top A x \geq 0, \forall x \in \mathbb{R}^n$. Finally, we denote by \mathbf{e}_i the i^{th} unit vector.

3.1 Key elements in Performance Estimation

This section introduces the core concepts of the Performance Estimation Problem (PEP) framework, originally proposed in [DT14] and further developed in [THG17c], through a simple illustrative example. Consider a single iteration of the Gradient Descent (GD) on the class $\mathcal{F}_{0,L}$ of L -smooth convex functions. Assume that the goal is to derive the worst-case upper bound on $f(x_1) - f(x_*)$, valid for all functions in $\mathcal{F}_{0,L}$, under the assumption that $\|x_0 - x_*\|^2 \leq 1$, and for all dimensions d of the problem. The tightest possible such bound can be formulated as the solution to the following optimization problem:

$$\begin{aligned}\max_{\substack{d \in \mathbb{N}, \\ f \in \mathcal{F}_{0,L}, \\ x_0, x_* \in \mathbb{R}^d}} \left\{ f(x_1) - f(x_*) \mid \|x_0 - x_*\|^2 \leq 1, \right. \\ \left. x_1 = x_0 - \alpha \nabla f(x_0), \|\nabla f(x_*)\|^2 = 0 \right\}.\end{aligned}\tag{3.1}$$

Finite-dimensional formulation of (3.1) As such, given the variable $f \in \mathcal{F}_{0,L}$, (3.1) is infinite-dimensional and untractable. However, as shown, e.g., in [DT14, THG17c], it can be equivalently reformulated as a maximization over finite-dimensional variables, under a set of constraints ensuring that the reformulation yields an optimal value matching that of (3.1). Typically, consider, as such variables, a set $S = \{(x_i, f_i, g_i)\}_{i=0,1,*} \in (\mathbb{R}^d \times \mathbb{R} \times \mathbb{R}^d)^3$, meant to encode the iterates of (GD) and evaluations of the function and (sub)gradient at these iterates. To obtain a finite-dimensional reformulation of (3.1) by maximizing over the variables in S , one should ensure (i) that S is consistent with the dynamics of the method and the problem assumptions, i.e., it satisfies (GD), $\|x_0 - x_*\|^2 \leq 1$, and $\|g_*\|^2 = 0$, and (ii)

that S indeed contains evaluations of a function in $\mathcal{F}_{\mu,L}$, i.e., that S is $\mathcal{F}_{0,L}$ -interpolable (Definition 2.6).

As illustrated in Chapter 2, for some function classes \mathcal{F} , including $\mathcal{F}_{0,L}$, the condition that a finite set S is \mathcal{F} -interpolable can be expressed via a set of algebraic conditions involving the elements of S . When such conditions are both necessary and sufficient for \mathcal{F} -interpolability, they are referred to as *interpolation conditions*, and they enable an exact reformulation of (3.1) as a finite-dimensional optimization problem. For $\mathcal{F}_{0,L}$, interpolation conditions are known, see Proposition 2.1. Hence, (3.1) can be equivalently reformulated as

$$\max_{d \in \mathbb{N}, S = \{(x_i, f_i, g_i)\}_{i=0,1,\star}} \left\{ f_1 - f_\star \mid S \text{ satisfies (2.15),} \right. \quad (3.2)$$

$$\left. \|x_0 - x_\star\|^2 \leq 1, x_1 = x_0 - \alpha g_0, \|g_\star\|^2 = 0 \right\}.$$

SDP expression of (3.1) While problem (3.2) is finite-dimensional, it remains difficult to solve due to its nonconvex nature, which stems from the scalar products between iterates and (sub)gradients that appear in the interpolation conditions (2.15).

To overcome this difficulty, [DT14, THG17c] proposed an alternative parametrization. Rather than optimizing over the set of variables S directly, they introduce a representation based on

- (i) the vector of function values $F = [f_\star, f_0, f_1]^\top$, (3.3)
- (ii) the *Gram matrix* $G = PP^\top$, where $P = [x_\star, x_0, g_0, g_1]^\top \in \mathbb{R}^{4 \times d}$.

By construction, G captures all inner products in x_i 's, g_i 's, $i \in \{0,1,\star\}$, including those involving x_1 and g_\star , which can be expressed via (GD) as an affine combination of the columns of P . Consequently, any quadratic quantity such as $\|x_1\|^2$ can be written as a linear function of the entries of G . Specifically, to access elements of G and F , we define selection vectors

- $\mathbf{g}_\star = \mathbf{0} \in \mathbb{R}^4$, $\mathbf{x}_\star = \mathbf{e}_1 \in \mathbb{R}^4$, $\mathbf{f}_\star = \mathbf{e}_1 \in \mathbb{R}^3$,
- For $k \in \{0,1\}$, $\mathbf{g}_k = \mathbf{e}_{k+3} \in \mathbb{R}^4$ and $\mathbf{f}_k = \mathbf{e}_{k+2} \in \mathbb{R}^3$,
- $\mathbf{x}_0 = \mathbf{e}_2 \in \mathbb{R}^4$, and $\mathbf{x}_1 = \mathbf{x}_0 + \alpha \mathbf{g}_0$.

These vectors allow accessing all f_i 's in F , and scalar products of x_i, g_i 's in

G. That is, $\forall i, j \in 0, 1, \star$,

$$\begin{aligned} \mathbf{x}_i^\top G \mathbf{x}_j &= \langle x_i, x_j \rangle, & \mathbf{g}_i^\top G \mathbf{g}_j &= \langle g_i, g_j \rangle, & \mathbf{x}_i^\top G \mathbf{g}_j &= \langle x_i, g_j \rangle, \\ F^\top \mathbf{f}_i &= f_i. \end{aligned}$$

This leads to an equivalent reformulation of (3.2) as a SDP:

$$\begin{aligned} \max_{\substack{F \in \mathbb{R}^3, \\ G \in \mathbb{S}_+^4}} \left\{ F^\top (\mathbf{f}_1 - \mathbf{f}_\star) \mid \forall i, j \in \{0, 1, \star\}, \right. \\ \left. F^\top (\mathbf{f}_j - \mathbf{f}_i) \geq \mathbf{g}_i^\top G (\mathbf{x}_j - \mathbf{x}_i) + \frac{1}{2L} (\mathbf{g}_i - \mathbf{g}_j)^\top G (\mathbf{g}_i - \mathbf{g}_j), \right. \\ \left. \text{and } (\mathbf{x}_0 - \mathbf{x}_\star)^\top G (\mathbf{x}_0 - \mathbf{x}_\star) \leq 1 \right\}, \end{aligned}$$

or, equivalently,

$$\max_{\substack{F \in \mathbb{R}^3, \\ G \in \mathbb{S}_+^4}} \left\{ F^\top b_{\text{obj}} \mid F^\top b_{ij} + \text{Tr}(A_{ij}G) \leq 0, \forall i, j \in \{0, 1, \star\}, \text{Tr}(A_0G) \leq 1 \right\}, \quad (3.4)$$

for some matrices $A_{ij}, A_0, A_{\varepsilon_0} \in \mathbb{S}^4$ and vectors $b_{\text{obj}}, b_{ij} \in \mathbb{R}^3$ that encode the objective and constraints of (3.2). The method dynamics (e.g., access to x_1) and the optimality condition are incorporated implicitly through the construction of these matrices and vectors.

Obtaining a dimension-dependent worst-case performance guarantee, e.g., an univariate worst-case guarantee, would require adding a rank constraint on the Gram matrix in (3.4), e.g., $\text{rank}(G) \leq 1$. This would render the problem non-convex.

Retrieving analytical proofs from SDP formulations Solving (3.4) yields a tight *numerical* performance guarantee for (GD) over $\mathcal{F}_{0,L}$ after a single iteration. However, one can extract an analytical proof from this numerical result.

As discussed in Chapter 1, most analytical proofs in first-order optimization consist of combining inequalities involving quantities accessible to the method (i.e., elements of S), which encode the assumptions on the problem class and the method's dynamics. That is, such inequalities are precisely the constraints appearing in (3.4). The dual coefficients associated with these constraints correspond to the optimal weights for which this

combination of inequalities yields the tightest possible guarantee. Thus, identifying analytical expressions for the dual coefficients in (3.4) (as functions of the method's parameters and the properties of the function class), enables one to convert the numerical PEP result into an analytical proof.

Any relaxation in this process, e.g., omitting some inequalities in the PEP formulation or approximating the dual variables using simpler expressions, may lead to suboptimal guarantees. Nevertheless, these relaxed bounds often remain significantly tighter than manually derived analytical guarantees. For more details, see, e.g., [GDT23].

Solving PEPs with *necessary* conditions A PEP formulation based on a non-exact algebraic representation of the problem's assumptions, such as using necessary but not sufficient conditions to characterize the function class, yields the best guarantee possible given the set of inequalities used in the formulation, but might be overconservative with respect to the initial function class. In other words, the resulting guarantee is optimal for a broader class of problems than the original one. Note that this broader class depends not only on the conditions used, but also on the problem solved, i.e., on the method, the performance measure, and the number of iterations analyzed.

3.2 Problem classes encompassed by the PEP framework

Although introduced here on a simple illustrative example, the PEP framework is broadly applicable and enables the derivation of tight worst-case guarantees for a wide variety of optimization problems. These guarantees are obtained by solving SDPs, provided that the ingredients of the analysis—namely, the algorithm, the function or operator class, the performance criterion, and the initial conditions—can be encoded (or approximated, e.g., via necessary interpolation conditions) in a *Gram-representable* way. Intuitively, this means that all expressions involved in the formulation of the PEP (e.g., objective and constraints) must depend linearly on function values and on scalar products between iterates and (sub)gradients.

Definition 3.1 (Gram-representable expression). Let $N \in \mathbb{N}$, and consider a finite set $S = \{(x_i, f_i, g_i)\}_{i=0,1,\dots,N, \star} \subset (\mathbb{R}^d \times \mathbb{R} \times \mathbb{R}^d)^{N+2}$. An algebraic expression p involving the elements of S is said to be *Gram-representable* if

it can be written using a finite number of linear operations involving only the function values f_j and scalar products of the form

$$\langle g_i, g_j \rangle, \quad \langle g_i, x_j \rangle, \quad \text{and} \quad \langle x_i, x_j \rangle, \quad \text{for all } i, j \in 0, 1, \dots, N, \star.$$

Given a Gram-representable expression p on $S = \{(x_i, f_i, g_i)\}_{i=0,1,\dots,N,\star}$, we equivalently say S satisfies p , or the Gram representation of S , i.e., the pair (F, G) as defined in (3.3), satisfies p .

Methods straightforwardly analyzable through SDP formulations encompasses a broad range of practical settings, including, e.g., (proximal/accelerated/inexact) gradient descent on all function classes introduced in Chapter 2, where the initial condition and performance measure may depend on $\|x_0 - x_\star\|^2$, $\|g_0\|^2$ for $g_0 \in \partial f(x_0)$, or $f(x_0) - f(x_\star)$, with analogous expressions involving x_N , g_N , and $f(x_N)$ in the performance criterion, see, e.g., [THG17a]. In addition, variational inequalities on classical operator classes also fit the PEP framework, see, e.g., [RTBG20].

Remark 3.1. Gram-representable expressions yield PEP formulations that are dimension-independent, and convex. PEP formulations can also be written with respect to non Gram-representable expressions, but solving such formulations may be less efficient.

3.3 A potential functions approach to performance analysis

This section briefly reviews an alternative to the PEP framework for the automated analysis of optimization algorithms, i.e, the automated construction of potential (or Lyapunov) functions.

This line of work was first introduced via the Integral Quadratic Constraints (IQC) framework [LRP16], which automatically computes the best *linear* convergence rate of a given method, i.e., the smallest ρ such that the error decreases like $\mathcal{O}(\rho^N)$ with N the number of iterations, by constructing quadratic Lyapunov functions that decrease over a single iteration. The framework was later extended to certify sublinear convergence rates in [HL17], i.e., smaller than linear rates. However, the resulting guarantees are inherently non-tight, as the IQC approach does not rely on interpolation conditions to describe function classes. In contrast, a methodology for certifying *exact* (sub)linear convergence using potential functions, merging the IQC and the PEP approach, has been developed in [TVSL18, TB19], and

further studied in, e.g., [ULG25, MTB23].

Given a pre-specified structure for a quadratic potential function, the methodology consists in optimizing the function's coefficients to yield the tightest possible convergence guarantee. For linear convergence, analyzing a single iteration is sufficient; for sublinear rates, multiple iterations must be considered simultaneously.

Compared to the PEP approach, this potential function-based method is more computationally and algebraically tractable. Moreover, it provides guarantees that hold for all time horizons, rather than for a fixed number of iterations N . However, it requires an initial guess for the potential function's structure and typically yields looser bounds than PEP for a given iteration count.

3.4 Other applications to computer-aided analysis

Computer-aided methods, aside from analyzing existing methods, can also be relied upon to (i) obtain lower bounds on the oracle complexity of first-order minimization of given problem classes [Dro17, DS20], and (ii) design new optimal (or efficient, given the level of relaxation introduced in the procedure) methods [BTB22, DT20, DGVPR24, KF16, KF21, Kim21, TD23, CHVSL18, VSFL17, DT16].

3.5 State of the art in computer-aided performance analysis

Since their introduction, computer-aided approaches to performance analysis have been used to analyze a wide range of optimization settings, which we briefly summarize. For further references, see, e.g., [Tay24, Bar21, Col24, Abb24].

The first problems analyzed using computer-aided techniques are deterministic first-order methods with fixed stepsizes (possibly time-varying), on Gram-representable problem classes. These include, e.g., the gradient method (and subgradient method) on various function classes [DT14, THG17c, Gri24, RGP24, AdKZ23a, TB19, GTD22, GEIGM22, ZG23], proximal gradient methods [THG18, KJH24], the difference-of-convex method [AdKZ24], the Alternating Direction Method of Multipliers [ZAdK24], Breg-

man gradient methods [DTdB22], the Chambolle-Pock method [BHG24], the heavy-ball method [GTD23], compressed gradient methods [TTD25], coordinate descent methods [SL17, AdKZ23b, KHG23, TB19], distributed methods [CH23], and variational and inclusion problems [RTBG20, LYR25, Lie21, GY20, GY24, GLG22, GTG22, GTHG23].

Beyond this setting, some extensions of computer-aided approaches include the analysis of:

- *Adaptive* first-order methods, by relaxing the conditions encoding the adaptive steps [DKGT17] or exactly analyzing the Polyak-stepsize scheme [BTd20].
- *Non-convex* PEP formulations, i.e., PEPs which do not admit a SDP reformulation, due, e.g., to constraints that are not Gram-representable, see, e.g., [RTBG20, DGFST24, DGVPR24].
- *Second-order* methods [DKGT20, Hil21]. Precisely, [DKGT20] used the PEP framework to analyze a single iteration of Newton’s method (NM) on the class of self-concordant functions, see, (7.2), as introduced in [NN94]. The analysis, however, is based a non-exact discrete representation of self-concordant functions, and is inherently restricted to the analysis of a single iteration as it relies on the use of a metric related to the starting iterate. Then, [Hil21, IH24] used tools from the IQC approach to tightly analyze the same setting. Extensions of computer-aided analyses to second-order optimization will be further discussed in Chapter 7.
- *Inexact* first-order methods, where evaluations of the gradient are corrupted by a deterministic perturbation, see, e.g., [DKGT20, Gan22, HSL21, VG24, THG17a, BTB22]. This straightforward extension simply requires adding the perturbations as additional variables to the PEP, which then computes the worst possible perturbations under given assumptions.
- *Stochastic* first-order methods as (SGD). In settings where the noise on the gradient follows a finite-support distribution, e.g., in the finite-sum and the block-coordinate settings, a first approach, yielding exact guarantees, consists in enumerating all scenarios and optimizing the average-case performance, as done, e.g., in [TB19, KHG23, CKPG25, HSR17]. Another approach consists in incorporating expectations of stochastic quantities, e.g., $\mathbb{E}[f(x_i)]$ for $i = 0, \dots, N$, directly

as variables in the associated SDPs, subject to *necessary* constraints imposed by the stochastic setting, see, e.g., [AdKZ23b, DS20]. Finally, [VSL21] analyzed stochastic methods by theoretically linking deterministic SDP solutions to robustness and convergence guarantees, and leveraging these SDPs to analyze or design methods. Extension of computer-aided analyses to stochastic optimization will be further discussed in Chapter 8.

Performance Estimation software The formulation and resolution of PEP formulations, under their SDP representation, is implemented in Matlab and Python toolboxes: PESTO [THG17b] and PEPit [GMG⁺22].

PART I
**An algebraic approach to
function and operator
interpolation**

4

A constraint-based approach to interpolation, and weakly convex optimization.

This chapter presents a *constraint-based* approach to interpolation. As discussed in Chapter 2, interpolation conditions are typically derived from given classes of functions or operators, as sets of constraints imposed on finite datasets to ensure their consistency with an element of the class. This perspective relies on the analytical properties of the function or operator classes.

The proposed constraint-based approach builds on the observation that, for most function classes, interpolation conditions characterize the class itself when imposed everywhere. This insight shifts the focus to the constraints p as the primary objects of study. Relying on the algebraic properties of p , we can determine: (i) whether the class they define when imposed everywhere is meaningful, in the sense that it encompasses non-trivial functions or operators, and (ii) whether these constraints constitute interpolation conditions for these classes, i.e., whether any function or operator satisfying them on a finite domain can be extended to the entire domain in a way that still satisfies p .

To assess this second property, we introduce the notion of *pointwise extensibility*. This concept captures whether a constraint permits consistent

extension to any additional point, rather than to the whole domain at once, and is therefore easier to verify than interpolability. Under suitable regularity conditions, the two notions turn out to be equivalent. Pointwise extensibility thus allows for an algebraic assessment of interpolability.

The chapter is structured as follows. Section 4.2 presents the constraint-based approach, and Section 4.3 formalizes the notion of pointwise extensibility. Then, Section 4.4 uses pointwise extensibility to derive counterexamples to the interpolability of classically used constraints, and Section 4.5 relies on this notion to obtain interpolation conditions for a class of non-convex non-smooth functions, that of weakly convex functions with bounded subgradient. Finally, Section 4.6 illustrates the constraint-based approach, by exhaustively analyzing a class of constraints and retaining those that are both meaningful and interpolable.

Notation This chapter considers non-convex functions, and specifically, weakly convex functions f , that is, functions such that, for some $\mu \geq 0$, $f + \frac{\mu}{2} \|\cdot\|^2$ is convex. For non-convex functions, the notion of a subgradient as introduced in (2.2) is not well defined. Hence, given a function $f : \mathbb{R}^d \rightarrow \mathbb{R}$, we introduce the notion of *regular* subgradient $v_x \in \mathbb{R}^d$ of f at x [RWW09, Chapter 8, Equation (3)], as vectors satisfying, for all $y \in \mathbb{R}^d$,

$$f(y) \geq f(x) + \langle v_x, y - x \rangle + \omega(\|x - y\|), \quad (4.1)$$

where the $\omega(t)$ notation refers to any term satisfying $\lim_{t \rightarrow 0} \frac{\omega(t)}{t} = 0$. The set of regular subgradients of f at x constitute the regular subdifferential of f at x . For f convex, the regular subdifferential reduces to the subdifferential [RWW09, Proposition 8.12]. By an abuse of notation, throughout this chapter, we denote the regular differential of f at x by $\partial f(x)$, and refer to its regular subgradients as its subgradients.

4.1 Codes.

Codes for reproducing and verifying technical parts of this chapter can be found at

https://github.com/AnneRubbens/Func_Interp_Code/tree/main.

4.2 A constraint-based approach to interpolation

This section introduces the constraint-based approach to interpolation by (i) defining, from given constraints, associated function classes, and (ii) presenting the properties of constraints we are interested in.

We consider function (resp. operator) classes characterized by properties that can be translated into a set of algebraic inequalities involving a finite number of points, and evaluations of a function f and its derivatives (resp. operator Q). Formally, let $\mathcal{X} \subseteq \mathbb{R}^d$ and

$$F : \mathcal{X} \rightrightarrows \mathbb{R}^D,$$

a possibly set-valued mapping. We consider constraints

$$p : (\mathbb{R}^d \times \mathbb{R}^D)^m \rightarrow \mathbb{R}^\ell, \quad \ell \in \mathbb{N},$$

and say F satisfies p at some tuple (x_1, \dots, x_m) in \mathcal{X}^m if

$$p((x_1, F_{x_1}), \dots, (x_m, F_{x_m})) \leq 0,$$

for all $F_{x_i} \in F(x_i)$, $i \in [m]$, where the inequality is to be taken elementwise for $\ell \geq 2$.

Mapping classes The proposed constraint-based approach to interpolation stems from the observation that any such algebraic constraint p defines a class of mappings over $\mathcal{X} \subseteq \mathbb{R}^d$, simply by requiring the constraint to hold at all m -tuples in \mathcal{X} .

Definition 4.1 (Mapping classes). Given a constraint $p : (\mathbb{R}^d \times \mathbb{R}^D)^m \rightarrow \mathbb{R}^\ell$ and $\mathcal{X} \subseteq \mathbb{R}^d$, a mapping class defined by p on \mathcal{X} consists in all mappings F satisfying p on \mathcal{X} , i.e.,

$$\mathcal{F}_p(\mathcal{X}) := \left\{ F : \mathcal{X} \rightrightarrows \mathbb{R}^D : p((x_1, F_{x_1}), \dots, (x_m, F_{x_m})) \leq 0, \forall F_{x_i} \in F(x_i), i \in [m], \right. \\ \left. \text{and for all tuples } (x_1, \dots, x_m) \text{ in } \mathcal{X}^m \right\}.$$

Example: L -smooth convex functions

Let $D = 1 + d$, $m = 2$, $\ell = 1$, and $\mathcal{X} = \mathbb{R}^d$. Consider single-valued mappings

$$F : \mathbb{R}^d \rightarrow \mathbb{R} \times \mathbb{R}^d : x \rightarrow F(x) = (f(x), g(x)),$$

and the pairwise constraint $p'_{0,L} : (\mathbb{R}^d \times \mathbb{R} \times \mathbb{R}^d) \rightarrow \mathbb{R}^2$:

$$\begin{aligned} p'_{0,L}((x, f(x), g(x)), (y, f(y), g(y))) &\leq 0 & (4.2) \\ \Leftrightarrow \begin{cases} f(y) \geq f(x) + \langle g(x), y - x \rangle \\ f(y) \leq f(x) + \langle g(x), y - x \rangle + \frac{L}{2} \|x - y\|^2. \end{cases} \end{aligned}$$

Then,

$$\mathcal{F}_{p'_{0,L}}(\mathbb{R}^d) = \left\{ (f, g) : \mathbb{R}^d \rightarrow \mathbb{R}^{1+d} \mid \begin{cases} f(y) \geq f(x) + \langle g(x), y - x \rangle \\ f(y) \leq f(x) + \langle g(x), y - x \rangle + \frac{L}{2} \|x - y\|^2 \end{cases} \forall x, y \in \mathbb{R}^d \right\},$$

is the class of L -smooth convex functions on \mathbb{R}^d , where $g(x) = \nabla f(x)$.

For constraints p characterizing differentiable function classes when imposed everywhere, e.g., (4.2), $\mathcal{F}_p(\mathbb{R}^d)$ exactly describes the classically defined associated function class, e.g., $\mathcal{F}_{p'_{0,L}}(\mathbb{R}^d)$ encompasses all convex L -smooth functions defined everywhere. Constraints p characterizing non-differentiable function classes when imposed everywhere, however, require the introduction of a technical subtlety to match classical function classes. Indeed, as such, $\mathcal{F}_p(\mathbb{R}^d)$ encompasses all functions satisfying p on the whole domain, for *at least one element in their subdifferential*, by contrast with classically defined function classes, requiring satisfaction of p for *all* elements in the subdifferential. To close the gap between $\mathcal{F}_p(\mathbb{R}^d)$ and classically defined function classes, we introduce the notion of *maximal* elements in $\mathcal{F}_p(\mathcal{X})$, as mappings $F \in \mathcal{F}_p(\mathcal{X})$ for which there exists no proper extension in $\mathcal{F}_p(\mathcal{X})$, and thus satisfy p for all subgradients in their subdifferential.

Definition 4.2 (Class of maximal elements in $\mathcal{F}_p(\mathcal{X})$). Let $p: (\mathbb{R}^d \times \mathbb{R}^D)^m \rightarrow \mathbb{R}^\ell$, and $\mathcal{X} \subseteq \mathbb{R}^d$. We let

$$\begin{aligned} \mathcal{F}_p^m(\mathcal{X}) := \{F: \mathcal{X} \rightrightarrows \mathbb{R}^D \mid F \in \mathcal{F}_p(\mathcal{X}), \\ \exists \bar{F} \in \mathcal{F}_p(\mathcal{X}) \text{ s.t. } \text{Graph}(F) \subseteq \text{Graph}(\bar{F})\}, \end{aligned} \quad (4.3)$$

where $\text{Graph } F := \{(x, F_x) \in \mathcal{X} \times \mathbb{R}^D \mid F_x \in F(x)\}$ is the graph of F .

We denote by $\mathcal{F}_p^m(\mathcal{X})$ the class of maximal functions in $\mathcal{F}_p(\mathcal{X})$. By Zorn's Lemma, see, e.g., [Rud87, Theorem 4.21], any mapping $F \in \mathcal{F}_p(\mathcal{X})$ admits an extension $\bar{F} \in \mathcal{F}_p^m(\mathcal{X})$, hence from an interpolation point of view considering one or another set is equivalent.

Proposition 4.1. Let $p: (\mathbb{R}^d \times \mathbb{R}^D)^m \rightarrow \mathbb{R}^\ell$, and $\mathcal{X} \subseteq \mathbb{R}^d$. Given any $F \in \mathcal{F}_p(\mathcal{X})$, there exists some $\bar{F} \in \mathcal{F}_p^m(\mathcal{X})$ such that

$$\text{Graph}(F) \subseteq \text{Graph}(\bar{F}).$$

Proof. The proof follows the same lines as that of [BC17, Theorem 20.21], treating maximally monotone operators. Let $\mathcal{P} = \{H \in \mathcal{F}_p(\mathcal{X}) \mid \text{Graph } \bar{F} \subseteq \text{Graph } H\}$, which is non-empty, along with the partial order on \mathcal{P} .

$$(H_1 \preceq H_2) \Leftrightarrow \text{Graph } H_1 \subseteq \text{Graph } H_2, \forall H_1, H_2 \in \mathcal{P}.$$

Given any chain \mathcal{C} in \mathcal{P} , let \bar{H} be defined as

$$\text{Graph } \bar{H} = \bigcup_{H \in \mathcal{C}} \text{Graph } H.$$

Then $\bar{H} \in \mathcal{P}$, and \bar{H} is an upper bound for \mathcal{C} . Hence, Zorn's lemma guarantees the existence of a maximal element $h \in \mathcal{P}$, which extends \bar{F} by definition of \mathcal{P} , and cannot be extended since maximal in \mathcal{P} . \square

Example: convex functions

Let $D = 1 + d$, $m = 2$, $\ell = 1$, and $\mathcal{X} = \mathbb{R}^d$. Consider set-valued mappings

$$F: \mathbb{R}^d \rightrightarrows \mathbb{R} \times \mathbb{R}^d: x \rightarrow F(x) = (f(x), g(x)),$$

and the pairwise constraint $p_{0,\infty} : (\mathbb{R}^d \times \mathbb{R} \times \mathbb{R}^d) \rightarrow \mathbb{R}$:

$$p((x, f(x), g_x), (y, f(y), g_y)) = -f(y) + f(x) + \langle g_x, y - x \rangle. \quad (p_{0,\infty})$$

Then,

$$\mathcal{F}_{p_{0,\infty}}^m(\mathbb{R}^d) = \{(f, g) : \mathbb{R}^d \rightarrow \mathbb{R}^{1+d} \mid \\ f(x) \geq f(y) + \langle g_x, y - x \rangle, \forall x, y \in \text{dom} f, \forall g_x \in g(x)\}$$

is the class of convex finite-valued, i.e., $\text{dom} f = \mathbb{R}^d$, functions on \mathbb{R}^d , where $g(x) = \partial f(x)$.

We now describe properties of constraints we are interested in, that is, they should be interpolation conditions, and the class they characterize when imposed everywhere should be meaningful.

Interpolation conditions. Let $\mathcal{X} \subset \bar{\mathcal{X}} \subseteq \mathbb{R}^d$, where $\bar{\mathcal{X}}$ is convex, and p a constraint. Trivially, any mapping \bar{F} satisfying p on $\bar{\mathcal{X}}$ can be restricted to a mapping F satisfying p on \mathcal{X} . However, the contrary is not necessarily true, as some mappings F satisfying p on \mathcal{X} may not admit an extension to $\bar{\mathcal{X}}$, in a way as to still satisfying p , see, e.g., the example of smooth convex functions introduced in Chapter 2, or in [THG17c, Figure 1]. We are interested in constraints p for which such an extension is always possible, referred to as *interpolable constraints*, or interpolation conditions for $\mathcal{F}_p(\bar{\mathcal{X}})$ (or $\mathcal{F}_p^m(\bar{\mathcal{X}})$). Interpolation conditions for a class \mathcal{F} were already defined in Definition 2.6, as necessary and sufficient conditions to impose on any finite set to ensure its consistency with an element in \mathcal{F} . However, we propose a slightly different definition, which considers interpolability as an algebraic property of a given constraint.

Definition 4.3 (Interpolable constraint). Let $\bar{\mathcal{X}} \subseteq \mathbb{R}^d$ convex. We say a constraint p is interpolable, or is an interpolation condition for $\mathcal{F}_p(\bar{\mathcal{X}})$ (resp. $\mathcal{F}_p^m(\bar{\mathcal{X}})$) if

$$\forall \text{ finite } \mathcal{X} \subset \bar{\mathcal{X}}, \forall F \in \mathcal{F}_p(\mathcal{X}), \exists \bar{F} \in \mathcal{F}_p(\bar{\mathcal{X}}) \text{ (resp. } \mathcal{F}_p^m(\bar{\mathcal{X}})) \\ \text{s.t. } F(x) = \bar{F}(x) \forall x \in \mathcal{X}.$$

If p satisfies Definition 2.6, then p is an interpolation constraint for the

class it defines, in the sense of Definition 4.3.

Proposition 4.2. *Let \mathcal{F} be a function or operator class, and p be an interpolation condition for \mathcal{F} , in the sense of Definition 2.6. Then p is an interpolation condition in the sense of Definition 4.3 for $\mathcal{F}_p(\mathbb{R}^d)$.*

Proof. First, we show $\mathcal{F} \subseteq \mathcal{F}_p(\mathbb{R}^d)$. Consider any $F \in \mathcal{F}$, and suppose that, at some m -tuple in \mathbb{R}^d , F does not satisfy p . Then, by interpolability of p , there exists no function in \mathcal{F} interpolating this tuple, contradicting $F \in \mathcal{F}$. Hence, $F \in \mathcal{F}_p(\mathbb{R}^d)$. In addition, consider any finite $\mathcal{X} = \{x_i\}_{i \in [N]} \subset \mathbb{R}^d$, and any $F \in \mathcal{F}_p(\mathcal{X})$. Since $S = \{(x_i, F_i)\}_{i \in [N]}$ satisfies p , it is interpolable by some $\bar{F} \in \mathcal{F} \subseteq \mathcal{F}_p(\mathbb{R}^d)$. Hence, there exists some $\bar{F} \in \mathcal{F}_p(\mathbb{R}^d)$ such that $F(x) = \bar{F}(x)$, $\forall x \in \mathcal{X}$, which concludes the proof. \square

The other direction does not hold: consider for instance the trivial constraint $p = 0$. This constraint is an interpolation constraint as in Definitions 4.3 and 2.6, but not necessarily for the same function classes. Indeed, p is an interpolation constraint as in Definition 4.3 for the class $\mathcal{F}_p(\mathbb{R}^d)$ containing all mappings, and an interpolation constraint as in Definition 2.6 for the class \mathcal{F} of continuous functions, since any finite dataset is consistent with a continuous function.

Function classes and differentially consistent constraints We are interested in constraints that, in addition to being interpolation constraints, define non-trivial mapping classes. In particular, suppose a constraint p is meant to describe a function class, by involving function values and first-order derivatives, e.g., (4.2) or $(p_{0,\infty})$. Mappings of interest then take the form $F : \mathcal{X} \rightrightarrows \mathbb{R}^{1+d} : x \rightrightarrows (f(x), g(x)) \in \mathbb{R} \times \mathbb{R}^d$. While meant to satisfy $g(x) \in \partial f(x)$, the mappings $g(x)$ and $f(x)$ are *a priori* unrelated, other than satisfying a certain algebraic relationship. We seek constraints whose satisfaction everywhere implies $g(x) \in \partial f(x)$, $\forall x \in \mathbb{R}^d$, referred to as *differentially consistent*.

Definition 4.4 (Differentially consistent constraint). A constraint $p : (\mathbb{R}^d \times \mathbb{R} \times \mathbb{R}^d) \rightarrow \mathbb{R}^\ell$ *differentially consistent* if, given any $F : \mathbb{R}^d \rightrightarrows \mathbb{R}^{1+d} : x \rightrightarrows (f(x), g(x))$, $F \in \mathcal{F}_p^m(\mathbb{R}^d)$ implies

$$g(x) = \partial f(x), \forall x \in \mathbb{R}^d.$$

For instance, given the definition (4.1) of a subgradient, a pairwise ($m = 2$) constraint p is differentially consistent if satisfaction of p at some pair

$((x_i, f_i, g_i), (x_j, f_j, g_j)) \in (\mathbb{R}^d \times \mathbb{R} \times \mathbb{R}^d)^2$ implies

$$f_j \geq f_i + \langle g_i, x_j - x_i \rangle + \omega(\|x_i - x_j\|). \quad (4.4)$$

Example: convex functions

Consider the constraint $(p_{0,\infty})$, characterizing convex functions. It holds that $(p_{0,\infty})$ is differentially consistent since its satisfaction at $((x_i, f_i, g_i), (x_j, f_j, g_j))$ implies $f_j \geq f_i + \langle g_i, x_j - x_i \rangle$, hence (4.4).

Example 4.1. The trivial pairwise constraint $p((x_i, f_i, g_i), (x_j, f_j, g_j)) = 0$ is not differentially consistent since satisfied by all mappings F on \mathbb{R}^d , e.g., by $f(x) = 0$, $g(x) = x, \forall x \in \mathbb{R}^d$, which do not satisfy $g(x) = \partial f(x)$.

Section 4.2 introduced the proposed algebraic approach to interpolability in a general context. For the sake of readability, we now introduce lighter notation related to the cases we analyze in what follows.

Compact notation for constraints characterizing function classes Consider a constraint $p : (\mathbb{R}^d \times \mathbb{R} \times \mathbb{R}^d)^m \rightarrow \mathbb{R}^\ell$ characterizing a first-order function class, i.e., via evaluations of function and (sub)gradient values, and some finite set $\mathcal{X} \subset \mathbb{R}^d$, along with a mapping $F = (f, g) \in \mathcal{F}_p(\mathcal{X})$. We denote by $S = \{(x_i, f_i, g_i)\}_{i \in [N]} \subset (\mathbb{R}^d \times \mathbb{R} \times \mathbb{R}^d)^N$ the combination of all points in \mathcal{X} and their associated values under F . That is, if g is multi-valued, S may include multiple entries with the same x_i , associated to different values of g_i .

When considering operator classes, mappings of interest take the form $F : \mathbb{R}^d \rightarrow \mathbb{R}^d : x \rightarrow Q(x)$, and in this case we denote by $S = \{(x_i, q_i)\}_{i \in [N]}$ the dataset to be extended, where again a multivalued mapping is accounted for by letting several points in S be equal, with different associated q_i 's. To avoid notation clashes, we refer to properties of operators using the letter q instead of p .

We denote with a superscript the points at which p is evaluated, e.g., for a pairwise constraint ($m = 2$),

$$p^{xy} = p((x, f_x, g_x), (y, f_y, g_y)), \text{ for some } (f_x, g_x) \in F(x), (f_y, g_y) \in F(y),$$

or

$$p^{ij} = p((x_i, f_i, g_i), (x_j, f_j, g_j)),$$

and we say S satisfies p if $p^{ij} \leq 0$ for all $i, j \in [N]$. Finally, we denote 2-wise by pairwise, 3-wise by terwise, etc.

Infinite-valued function classes Throughout, to avoid additional technical difficulties, we restrict attention to finite-valued functions. In principle, one could explicitly allow function classes whose elements are defined on \mathbb{R}^d , but take finite values only on an *effective domain* $\bar{\mathcal{X}} \subseteq \mathbb{R}^d$.

Suppose first this effective domain to be predetermined, e.g., consider the class of functions defined everywhere, but finite-valued and satisfying a constraint p on a given set $\bar{\mathcal{X}}$ only. Definition 4.3 enables the analysis of such functions straightforwardly, since extending elements in $\mathcal{F}_p(\bar{\mathcal{X}})$ to functions satisfying p on their effective domain $\bar{\mathcal{X}}$, and infinite outside of $\bar{\mathcal{X}}$ is straightforward. To tackle this case, it thus suffices to slightly modify Definition 4.1 to account for this notion of effective domain.

On the other hand, classical function classes as the class of convex functions, require the *existence* of such an effective domain rather than predefining it. A first approach to tackle this case would be to slightly modify Definition 4.3, so as to allow $\bar{\mathcal{X}}$ to depend on \mathcal{X} , e.g., to be its convex hull [RWW09, Chapter 2.E]. A second approach matching these classical classes would be to introduce the mapping family

$$\mathcal{F}_p^{m,\infty}(\mathbb{R}^d) := \left\{ F : \mathbb{R}^d \rightrightarrows \bar{\mathbb{R}}^D \cup \emptyset \mid \exists \bar{\mathcal{X}} \subseteq \mathbb{R}^d \text{ convex s.t. } F|_{\bar{\mathcal{X}}} \in \mathcal{F}_p^m(\bar{\mathcal{X}}) \text{ and } F(x) = (+\infty, \emptyset), x \in \mathbb{R}^d \setminus \bar{\mathcal{X}} \right\}, \quad (4.5)$$

where $F|_{\bar{\mathcal{X}}}$ denotes the restriction of F to $\bar{\mathcal{X}}$, and requiring extension to $\mathcal{F}_p^{m,\infty}(\mathbb{R}^d)$ instead of $\mathcal{F}_p^m(\mathbb{R}^d)$.

4.3 Pointwise extensibility

Based on Definition 4.3, a constraint p is an interpolation constraint for $\mathcal{F}_p(\mathbb{R}^d)$ if any mapping satisfying p on a finite domain admits an extension satisfying p on \mathbb{R}^d . We now introduce the weaker notion of a *pointwise*

extensible constraint, based on the idea of extending a mapping satisfying a constraint p on a finite set to just one (arbitrary) point, in the spirit of what is done in, e.g., [AM17a, Gru09] with respect to potentially infinite sets. Verifying pointwise extensibility of a constraint can be done algebraically, and we show in Theorem 4.1, under regularity assumptions, that pointwise extensible and interpolation constraints are equivalent.

Definition 4.5 (Pointwise extensible constraint). Let $\bar{\mathcal{X}} \subseteq \mathbb{R}^d$ convex. We say a constraint p is pointwise extensible on $\bar{\mathcal{X}}$ if

$$\begin{aligned} \forall \text{ finite } \mathcal{X} \subset \mathbb{R}^d, \forall z \in \bar{\mathcal{X}}, \forall F \in \mathcal{F}_p(\mathcal{X}), \exists \bar{F} \in \mathcal{F}_p(\mathcal{X} \cup z) \\ \text{s.t. } F(x) = \bar{F}(x) \quad \forall x \in \mathcal{X}. \end{aligned}$$

Example: convex functions

Consider the constraint $(p_{0,\infty})$, characterizing convex functions. For some $N \in \mathbb{N}$, consider $S = \{(x_i, f_i, g_i)\}_{i \in [N]}$, where S satisfies $p_{0,\infty}$. Given any $z \in \mathbb{R}^d$, let

$$i_\star = \operatorname{argmax}_{i \in [N]} f_i + \langle g_i, x - x_i \rangle, \quad f_z = f_{i_\star} + \langle g_{i_\star}, x - x_{i_\star} \rangle \text{ and } g_z = g_{i_\star}.$$

By definition of i_\star , $f_z \geq f_i + \langle g_i, x - x_i \rangle \quad \forall i \in [N]$. In addition, since $p_{0,\infty}^{j_{i_\star}} \geq 0$,

$$f_j \geq f_{i_\star} + \langle g_{i_\star}, x_j - x_{i_\star} \rangle = f_z + \langle g_z, x_j - x \rangle \quad \forall j \in [N].$$

Hence, these f_z and g_z extend $p_{0,\infty}$ to z and $p_{0,\infty}$ is pointwise extensible on \mathbb{R}^d .

Under some regularity assumptions, pointwise extensible and interpolable constraints are equivalent. By definition, the first implication holds.

Proposition 4.3 (From interpolability to pointwise extensibility). Let $\bar{\mathcal{X}} \subseteq \mathbb{R}^d$ convex, and a constraint p . If p is an interpolation constraint for $\mathcal{F}_p(\bar{\mathcal{X}})$, (or $\mathcal{F}_p^m(\bar{\mathcal{X}})$), then p is pointwise extensible on $\bar{\mathcal{X}}$.

Proposition (4.3) allows deriving proofs of non-interpolability of a constraint p , see Section 4.4 for an illustration. The converse of Proposition 4.3 holds under additional technical conditions. Indeed, pointwise extensibility allows iteratively extending any $F \in \mathcal{F}_p(\mathcal{X})$ to the rationals, i.e., to

some $\bar{F} \in \mathcal{F}_p(\mathbb{Q}^d)$. Then, one can extend \bar{F} to \mathbb{R}^d , taking its limit on \mathbb{Q}^d , provided (i) such a limit exists, and (ii) p is continuous, hence still satisfied at the limit. We formalize these requirements by introducing the notion of regular extensibility, before providing sufficient conditions for regular extensibility of a constraint describing first-order function classes.

Definition 4.6 (Regularly extensible constraint).

Let $\bar{\mathcal{X}} \subseteq \mathbb{R}^d$ convex. We say a constraint p is *regularly extensible* on $\bar{\mathcal{X}}$ if

1. p is continuous.
2. Given a finite set $\mathcal{X} \subseteq \bar{\mathcal{X}}$, $\mathbb{Q}^d \cap \bar{\mathcal{X}}$ the set of rationals in $\bar{\mathcal{X}}$, and $F \in \mathcal{F}_p(\mathcal{X})$,

$$\exists \bar{F} \in \mathcal{F}_p((\mathbb{Q}^d \cap \bar{\mathcal{X}}) \cup \mathcal{X}) : F(x) = \bar{F}(x), x \in \mathcal{X},$$

such that alongside any Cauchy sequence $(x_k) \subset \mathbb{Q}^d \cap \bar{\mathcal{X}}$, there exists a convergent subsequence in $(\bar{F}(x_k))$.

Lemma 4.1 (Sufficient conditions for regular extensibility). *Let $\bar{\mathcal{X}} \subseteq \mathbb{R}^d$ convex, and $p : (\mathbb{R}^d \times \mathbb{R} \times \mathbb{R}^d)^2 \rightarrow \mathbb{R}^\ell$ a pairwise continuous, differentially consistent, pointwise extensible constraint. If any of the following conditions hold:*

- Given $\{(x_i, f_i, g_i), (x_i, f_i, g_i)\} \in (\mathbb{R}^d \times \mathbb{R} \times \mathbb{R}^d)^2$, satisfaction of p^{ij} and p^{ji} implies

$$\|g_i - g_j\| \leq \omega(\|x_i - x_j\|); \quad (4.6)$$

- Given $\{(x_i, f_i, g_i), (x_i, f_i, g_i)\} \in (\mathbb{R}^d \times \mathbb{R} \times \mathbb{R}^d)^2$, satisfaction of p^{ij} and p^{ji} implies

$$\|g_i\| \leq A\|g_j\| + B, \text{ where } A, B < +\infty, \text{ or}; \quad (4.7)$$

- Given any finite $\mathcal{X} \subset \bar{\mathcal{X}}$ and $z \in \bar{\mathcal{X}}$, $\forall F = (f, g) \in \mathcal{F}_p(\mathcal{X})$,

$$\exists \bar{F} \in \mathcal{F}_p(\mathcal{X} \cup z) \quad (4.8)$$

$$\text{s.t. } \begin{cases} F(x) = \bar{F}(x) \quad \forall x \in \mathcal{X} \\ \bar{F}(z) = (\bar{f}(z), \bar{g}(z)) \text{ satisfies } \|\bar{g}(z)\| \leq \max_{x_i \in \mathcal{X}, g_i \in g(x_i)} \|g_i\|; \end{cases}$$

then p is regularly extensible on $\bar{\mathcal{X}}$.

Proof. We show existence of some $\bar{F} = (\bar{f}, \bar{g}) \in \mathcal{F}_p((\mathbb{Q}^d \cap \bar{\mathcal{X}}) \cup \mathcal{X})$: $F(x) = \bar{F}(x)$, $x \in \mathcal{X}$, such that alongside any Cauchy sequence $(x_k) \subset \mathbb{Q}^d \cap \bar{\mathcal{X}}$, it holds that (i) $(\bar{g}(x_k))$ is bounded, hence contains a convergent subsequence [BS00, 3.4.8], and (ii) $(\bar{f}(x_k))$ is a Cauchy sequence, hence a convergent sequence [BS00, 3.5.5].

Suppose first (4.6) or (4.7) is satisfied. By countability of $\mathbb{Q}^d \cap \bar{\mathcal{X}}$ and pointwise extensibility of p , F can be iteratively extended to $\mathbb{Q}^d \cap \bar{\mathcal{X}} \cup \mathcal{X}$ in a way as to satisfy p . Denote by $\bar{F} \in \mathcal{F}_p((\mathbb{Q}^d \cap \bar{\mathcal{X}}) \cup \mathcal{X})$ this extension of F , and consider any Cauchy sequence $(x_k) \subset \mathbb{Q}^d \cap \bar{\mathcal{X}}$. Suppose (4.6) is satisfied. Alongside (x_k) , since Lipschitz continuity preserves Cauchy sequences, see, e.g., [BS00, Theorem 5.4.7], $\bar{g}(x_k)$ is a Cauchy sequence, hence bounded [BS00, 3.4.8]. On the other hand, if (4.7) is satisfied, then

$$\|\bar{g}(x_k)\| \leq \min_{x_i \in \mathcal{X}, g_i \in g(x_i)} A \|g_i\| + B < +\infty, \forall k,$$

hence $(\bar{g}(x_k))$ is bounded by some $M < +\infty$.

Suppose now p satisfies (4.8). By countability of \mathbb{Q}^d and pointwise extensibility of p , F can be iteratively extended to $(\mathbb{Q}^d \cap \bar{\mathcal{X}}) \cup \mathcal{X}$ in a way as to satisfy p , and $\bar{g}(x) \leq \max_{x_i \in \mathcal{X}, g_i \in g(x_i)} \|g_i\|$, at all $x \in \mathbb{Q}^d \cap \bar{\mathcal{X}}$. Denote by $\bar{F} \in \mathcal{F}_p(\mathbb{Q}^d \cap \bar{\mathcal{X}} \cup \mathcal{X})$ this extension of F , and consider any Cauchy sequence $(x_k) \subset \mathbb{Q}^d$. Alongside this sequence,

$$\|\bar{g}(x_k)\| \leq \max_{x_i \in \mathcal{X}, g_i \in g(x_i)} \|g_i\| \forall k,$$

hence $(\bar{g}(x_k))$ is bounded by some $M < +\infty$.

We now show that $(\bar{f}(x_k))$ is a Cauchy sequence. Indeed, since p is differentially consistent, satisfaction of p at any pair $(y_i, y_j) \in \bar{\mathcal{X}}$ implies satisfaction of (4.4) at (y_i, y_j) and (y_j, y_i) , i.e.,

$$\begin{aligned} & \begin{cases} \bar{f}(y_i) - \bar{f}(y_j) & \geq \langle \bar{g}(y_j), y_i - y_j \rangle + \omega(\|y_i - y_j\|) \\ \bar{f}(y_i) - \bar{f}(y_j) & \leq \langle \bar{g}(y_i), y_i - y_j \rangle - \omega(\|y_i - y_j\|) \end{cases} \\ \Leftrightarrow & \begin{cases} \bar{f}(y_i) - \bar{f}(y_j) & \geq -\|\bar{g}(y_j)\| \|y_i - y_j\| + \omega(\|y_i - y_j\|) \\ \bar{f}(y_i) - \bar{f}(y_j) & \leq \|\bar{g}(y_i)\| \|y_i - y_j\| - \omega(\|y_i - y_j\|) \end{cases} \\ \Leftrightarrow & |\bar{f}(y_i) - \bar{f}(y_j)| \leq M \|y_i - y_j\| + |\omega(\|y_i - y_j\|)|. \end{aligned}$$

By definition of $\omega(\cdot)$, $\forall \varepsilon > 0, \exists \delta_1 \geq 0 : |t| \leq \delta_1 \Rightarrow |\omega(|t|)| \leq \varepsilon |t|$. Hence, for

any $\varepsilon > 0$, let $\delta = \min\{\frac{\varepsilon}{M+\varepsilon}, \delta_1\}$. Then,

$$\text{If } \|y_i - y_j\| \leq \delta, \text{ then } |\bar{f}(y_i) - \bar{f}(y_j)| \underset{\delta \leq \delta_1}{\leq} (M + \varepsilon)\delta \underset{\delta \leq \frac{\varepsilon}{M+\varepsilon}}{\leq} \varepsilon.$$

Hence, since $(y_j)_{j=1,2,\dots}$ is a Cauchy sequence, so is $(\bar{f}(y_j))_{j=1,2,\dots}$ \square

For instance, by Example 4.3, the characterization $(p_{0,\infty})$ of convex functions is regularly extensible. We now show that regular pointwise extensible constraints are interpolable.

Theorem 4.1 (From pointwise extensibility to interpolability)

Let $\bar{\mathcal{X}} \subseteq \mathbb{R}^d$ convex. If p is a regular pointwise extensible constraint on $\bar{\mathcal{X}}$, then p is an interpolation constraint for $\mathcal{F}_p(\bar{\mathcal{X}})$ and $\mathcal{F}_p^m(\bar{\mathcal{X}})$.

Proof. Let $N \in \mathbb{N}$, $\mathcal{X} = \{x_i\}_{i \in [N]} \subset \mathbb{R}^d$ and $F \in \mathcal{F}_p(\mathcal{X})$. We construct $\bar{F} \in \mathcal{F}_p(\mathbb{R}^d)$ satisfying $F(x_i) = \bar{F}(x_i)$, $\forall i \in [N]$. For every $x \in \mathcal{X}$, set $\bar{F}(x) = F(x)$. Then, we extend \bar{F} to the rationals in $\bar{\mathcal{X}}$ as follows. By regular pointwise extensibility of p , there exists a mapping $\bar{F} \in \mathcal{F}_p((\mathbb{Q}^d \cap \bar{\mathcal{X}}) \cup \mathcal{X})$, satisfying $F(x_i) = \bar{F}(x_i)$, $\forall i \in [N]$, and containing a convergence subsequence alongside any Cauchy sequence (y_j) . For each remaining $x \in \bar{\mathcal{X}} \setminus \mathbb{Q}^d \cap \bar{\mathcal{X}}$, since $\mathbb{Q}^d \cap \bar{\mathcal{X}}$ is dense in $\bar{\mathcal{X}}$, there exists a sequence of rationals (y_j) converging to x . Alongside this sequence, consider $\bar{F}(y_j)$, containing a subsequence converging to some \bar{F} , by regular extensibility of p . Set $\bar{F}(x) = \bar{F}$. By continuity of p , $\bar{F}(x)$ extends p to x . Hence, F can be extended to some $\bar{F}(x) \in \mathcal{F}_p(\bar{\mathcal{X}})$. Extension to $\mathcal{F}_p^m(\bar{\mathcal{X}})$ follows from Proposition 4.1. \square

We now use the notions introduced in Sections 4.2 and 4.3, to algebraically assess interpolability of given constraints, both to obtain counterexamples to their interpolability, or establish their interpolability.

4.4 Proving non-interpolability of a constraint

This section builds on pointwise extensibility (Definition 4.5), to develop a heuristic computing counterexamples to a constraint's interpolability, when such counterexamples exist.

Let $\bar{\mathcal{X}} \subseteq \mathbb{R}^d$ convex. By Proposition 4.3, to obtain a counterexample to the interpolability of a constraint p for $\mathcal{F}_p(\bar{\mathcal{X}})$, it suffices to exhibit some $\mathcal{X} \subset \mathbb{R}^d$, and $F \in \mathcal{F}_p(\mathcal{X})$, together with a point $z \in \bar{\mathcal{X}}$ at which p cannot be extended. Given F and z , verifying whether such an extension is possible reduces to solving an optimization problem.

Lemma 4.2 (Counterexamples to interpolability). *Let $\bar{\mathcal{X}} \subset \mathbb{R}^d$ convex, and $\mathcal{X} \subset \bar{\mathcal{X}}$ finite. Consider a constraint $p : (\mathbb{R}^d \times \mathbb{R}^D)^m \rightarrow \mathbb{R}^\ell$, $F \in \mathcal{F}_p(\mathcal{X})$, and $z \in \bar{\mathcal{X}}$. Let*

$$\tau_\star := \min_{\substack{\tau \in \mathbb{R}, \\ F_z \in \mathbb{R}^D}} \left\{ \tau \text{ s.t. } p(\pi((x_1, F_{x_1}), \dots, (x_{m-1}, F_{x_{m-1}}), (z, F_z))) \leq \tau, \quad (4.9) \right.$$

$$\left. \begin{aligned} & \forall (x_1, x_2, \dots, x_{m-1}) \in \mathcal{X}^{m-1}, \forall F_{x_i} \in F(x_i), i \in [m-1], \\ & \text{and for all permutation } \pi \end{aligned} \right\}.$$

Then, F admits an extension satisfying p at z if and only if $\tau_\star \leq 0$.

Proof. Constraint p can be extended at z if there exists some F_z such that

$$\bar{F} : \mathcal{X} \cup z \Rightarrow \mathbb{R}^D : \begin{cases} \bar{F}(x) = F(x) & x \in \mathcal{X}, \\ \bar{F}(z) = F_z, \end{cases}$$

belongs to $\mathcal{F}_p(\mathcal{X} \cup z)$, i.e., such that

$$p((x_1, \bar{F}_{x_1}), \dots, (x_m, \bar{F}_{x_m})) \leq 0 \quad \forall \text{ tuple in } (\mathcal{X} \cup z)^m, \forall \bar{F}_{x_i} \in \bar{F}(x_i), i \in [m].$$

Problem (4.9) computes the smallest relaxation τ needed to satisfy all these inequalities. If $\tau_\star \leq 0$, an extension exists. \square

Example: convex functions

Consider the constraint $(p_{0,\infty})$, characterizing convex functions, $S = \{(x_i, f_i, g_i)\}_{i=1,2}$ satisfying $(p_{0,\infty})$, and $\bar{\mathcal{X}} = \mathbb{R}^d$. Then, (4.9) becomes

$$\tau_\star := \min_{\substack{\tau \in \mathbb{R}, \\ f_z \in \mathbb{R}, g_z \in \mathbb{R}^d}} \left\{ \tau \text{ s.t. } \begin{aligned} & -f_z + f_i + \langle g_i, z - x_i \rangle \leq \tau, \quad i = 1, 2 \\ & -f_i + f_z + \langle g_z, x_i - z \rangle \leq \tau, \quad i = 1, 2 \end{aligned} \right\}.$$

By pointwise extensibility of $(p_{0,\infty})$, $\tau_* \leq 0$ for all S .

Whenever τ_* , the solution to (4.9), is positive, F and z provide a counterexample to the interpolability of p . Thus, for a given cardinality K , one can search for such a counterexample by identifying the "hardest" z , $\mathcal{X} \subset \bar{\mathcal{X}}$ of cardinality K and $F \in \mathcal{F}_p(\mathcal{X})$, i.e., those maximizing τ_* . In particular, let

$$\tau_*(K) := \max_{\substack{z \in \bar{\mathcal{X}} \\ \mathcal{X} \subset \bar{\mathcal{X}}: |\mathcal{X}|=K \\ F \in \mathcal{F}_p(\mathcal{X})}} \min_{\substack{\tau \in \mathbb{R}, \\ F_z \in \mathbb{R}^D}} \tau \quad (4.10)$$

$$\begin{aligned} \text{s.t. } & p(\pi((x_1, F_{x_1}), \dots, (x_{m-1}, F_{x_{m-1}}), (z, F_z))) \leq \tau, \quad (4.11) \\ & \forall (x_1, x_2, \dots, x_{m-1}) \in \mathcal{X}^{m-1}, \\ & \forall F_{x_i} \in F(x_i), i \in [m-1], \text{ and for all permutation } \pi. \end{aligned}$$

Then, p is a pointwise extensible constraint if and only if $\tau_*(K) \leq 0, \forall K \in \mathbb{N}$.

Lemma 4.2, in a second step, can also be relied upon to *strengthen* constraints p that are not interpolable, that is to obtain a constraint \tilde{p} equivalent to p when imposed everywhere, but satisfied by fewer mappings than p on subsets of \mathbb{R}^d , see Chapter 5. In general, (4.10) is not directly solvable due to its non-convexity and the combinatorial nature of the search over \mathcal{X} and $F \in \mathcal{F}_p(\mathcal{X})$. To address this, we use a heuristic that iteratively constructs such \mathcal{X} and F maximizing the solution to (4.9), see the code referenced in Section 4.1 for implementation details. This heuristic, which does not solve (4.10) exactly, does not provide guarantees of interpolability in the absence of a counterexample. However, as illustrated in Table 4.1, it effectively obtained such counterexamples for several classical characterizations of function classes that have sparked interest in the optimization community, that is,

- Weakly convex functions with bounded subgradients [Nur74, Nur73], with minimum f_* , classically characterized, for some $\mu, B \geq 0$, by

$$p_{-\mu, B}^{x, y} \leq 0 \Leftrightarrow \begin{cases} f(y) \geq f(x) + \langle g_x, y - x \rangle - \frac{\mu}{2} \|x - y\|^2, \\ \|g_x\| \leq B, \\ f(x) \geq f_*, \end{cases} \quad (p_{-\mu, B})$$

that holds $\forall x, y \in \mathbb{R}^d$, and $\forall g_x \in \partial f(x)$. This function class is mo-

tivated in Section 4.5. In addition, interpolation conditions for this class are derived in Theorem 4.2.

- Smooth functions satisfying a Łojasiewicz condition [Loj63, BDL07], with minimum f_* , classically characterized, for some $0 \leq \mu \leq L$, by

$$p'_{L,\mu,L}{}^{x,y} \leq 0 \Leftrightarrow \begin{cases} f(x) \leq f(y) + \langle \nabla f(y), x - y \rangle + \frac{L}{2} \|x - y\|^2, \\ f(x) \leq f_* + \frac{\|g(x)\|^2}{2\mu}, \\ f(x) \geq f_*, \end{cases} \quad (p'_{L,\mu,L})$$

that holds $\forall x, y \in \mathbb{R}^d$. This function class is motivated in Section 5.5.1, which also derives a strengthening of $(p'_{L,\mu,L})$.

- Uniformly convex functions [Cla36, JN14a], classically characterized, for some $q \geq 2$ and $\mu \geq 0$, by

$$p_{\mu,q}^{xy} \leq 0 \Leftrightarrow f(x) \geq f(y) + \langle g_y, x - y \rangle + \frac{\mu}{q} \|x - y\|^q, \quad (p_{\mu,q})$$

that holds $\forall x, y \in \mathbb{R}^d$, and $\forall g_x \in \partial f(x)$. This function class is motivated in Section 5.5.4, which also derives a strengthening of $(p_{\mu,q})$.

- Hölder smooth convex functions [Höl86, NY83], classically characterized, for some $\alpha \in (0, 1)$ and $L \geq 0$, by

$$p_{L,\alpha}^{xy} \leq 0 \Leftrightarrow f(x) \leq f(y) + \langle \nabla f(y), x - y \rangle + \frac{L \|x - y\|^{\alpha+1}}{\alpha + 1}, \quad (p_{L,\alpha})$$

that holds $\forall x, y \in \mathbb{R}^d$.

- Functions with M -Lipschitz continuous Hessian, characterized among others via a cubic bound

$$|f(x) - f(y) - \langle \nabla f(y), x - y \rangle - \frac{1}{2}(x - y)^T \nabla^2 f(y)(x - y)| \leq \frac{M}{6} \|x - y\|^3, \quad (p_{\nabla^2, M})$$

that holds $\forall x, y \in \mathbb{R}^d$. This function class is motivated in Chapter 7, which also derives interpolation conditions for this class in the univariate case, see Theorem 7.4.

| Function class | Counterexample |
|---|--|
| Weakly convex functions with bounded subgradient | $S = \left\{ \begin{pmatrix} 0 \\ 0 \\ B \end{pmatrix}, \begin{pmatrix} \frac{2B}{\mu} + 1 \\ \frac{\mu}{2} - 2B \\ B \end{pmatrix} \right\}, \quad z = \frac{2B}{\mu}$ |
| Smooth functions satisfying a Łojasiewicz condition | $S = \left\{ \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 \\ \frac{L}{2} \\ L \end{pmatrix} \right\}, \mu = \frac{L}{2}, \quad z = 0.5$ |
| Uniformly convex functions | $S = \left\{ \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 2 \\ \frac{\mu 2^p}{p} \\ \frac{\mu 2^p}{p} \end{pmatrix}, \begin{pmatrix} 1 \\ \frac{\mu(2^{p-1}-1)}{p} \\ \frac{\mu 2^{p-1}}{p} \end{pmatrix} \right\}, \quad z = 0.1$ |
| Holder smooth functions | $S = \left\{ \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 2 \\ L(1 + \frac{1}{\alpha})^\alpha \\ L(1 + \frac{1}{\alpha})^\alpha \end{pmatrix}, \begin{pmatrix} 1 \\ \frac{L}{2}(1 + \frac{1}{\alpha})^\alpha - \frac{L}{\alpha+1} \\ \frac{L}{2}(1 + \frac{1}{\alpha})^\alpha \end{pmatrix} \right\},$ $z = 0.5$ |
| Functions with Lipschitz Hessian | $S = \left\{ \begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 \\ -\frac{M}{6} \\ 0 \\ 0 \end{pmatrix} \right\}, \quad z = 0.5$ |

Table 4.1 Function classes whose classical descriptions are not interpolable, and corresponding counterexamples. In the counterexamples, $S = \{(x_i, f_i, g_i)\}$ for first-order functions, $S = \{(x_i, f_i, g_i, h_i)\}$ for functions with Lipschitz Hessians, where h_i accounts for the second-order derivative of f , and z is the new point at which extending p is infeasible. One can verify these counterexamples using the code referenced in Section 4.1.

4.5 Weakly convex optimization

This section considers the class $\mathcal{F}_{p-\mu, B}^m(\mathbb{R}^d)$ of weakly convex functions with bounded subgradients. After motivating this function class, we derive its interpolation conditions by relying (i) on a graphical intuition to

propose a candidate interpolation constraint for $\mathcal{F}_{p-\mu,B}^m(\mathbb{R}^d)$, and (ii) on pointwise extensibility to prove interpolability of this candidate. We then use these conditions to analyze the subgradient method (4.30) on $\mathcal{F}_{p-\mu,B}^m(\mathbb{R}^d)$.

4.5.1 Motivation

While convex and/or L -smooth functions offer appealing theoretical guarantees, many real-world problems fall outside these classes, motivating the study of broader properties. In the non-smooth non-convex setting, the class of weakly convex functions emerges as particularly compelling, as it balances the favorable properties of convex functions with broader applicability. This class includes, for example, any composition

$$f(x) := h(c(x)),$$

where h is convex and Lipschitz continuous and c is smooth [DP19, Lemma 4.2]. Such compositions are widespread in applications; see, e.g., [DG19, DD19] for formulations of robust principal component analysis, minimization of conditional value-at-risk, robust phase retrieval, and covariance matrix estimation. On the other hand, this class possesses convergence guarantees (to stationary points only) and a continuous measure of stationarity [Bur85, DP19], detailed in Section 4.5.5. Boundedness of the subgradients is added to weak convexity to guarantee the convergence of common algorithms such as the subgradient method, which may diverge on weakly convex functions alone, see, e.g., [BW21, DD19, DGVPR24, MJ20].

In addition, the subgradient method (4.30) on $\mathcal{F}_{p-\mu,B}^m(\mathbb{R}^d)$ has been analyzed both classically, see, e.g., [DD19], and in a computer-aided way [DGVPR24], using the Gram-representable constraint $(p-\mu,B)$. As highlighted in Section 4.4, this constraint is however not interpolable, hence, all results are a priori non-tight. This motivates the derivation of interpolation conditions for $\mathcal{F}_{p-\mu,B}^m(\mathbb{R}^d)$.

Remark 4.1. Interpolation conditions for $\mathcal{F}_{p-\mu,B}^m(\mathbb{R}^d)$ can be retrieved from interpolation conditions for the class of convex functions f satisfying

$$\|g_x - \mu x\| \leq B, \forall x \in \mathbb{R}^d, \forall g_x \in \partial f(x).$$

However, since we only resort to algebraic tools, analyzing either formulation is equivalent. Hence, we consider $\mathcal{F}_{p-\mu,B}^m(\mathbb{R}^d)$ for straightforward comparison with [DDMP18, DGVPR24].

4.5.2 Strengthening of $(p_{-\mu,B})$

Motivated by a graphical, one-dimensional intuition, we derive a candidate interpolation constraint for $\mathcal{F}_{p_{-\mu,B}}^m(\mathbb{R}^d)$. Let $d = \mu = B = 1$, and $S = \{(0, 0, 1), (3, -\frac{3}{2}, 1)\}$, satisfying $(p_{-\mu,B})$. As illustrated in Figure 4.1, S cannot be interpolated by any function in $\mathcal{F}_{p_{-\mu,B}}^m(\mathbb{R}^d)$.

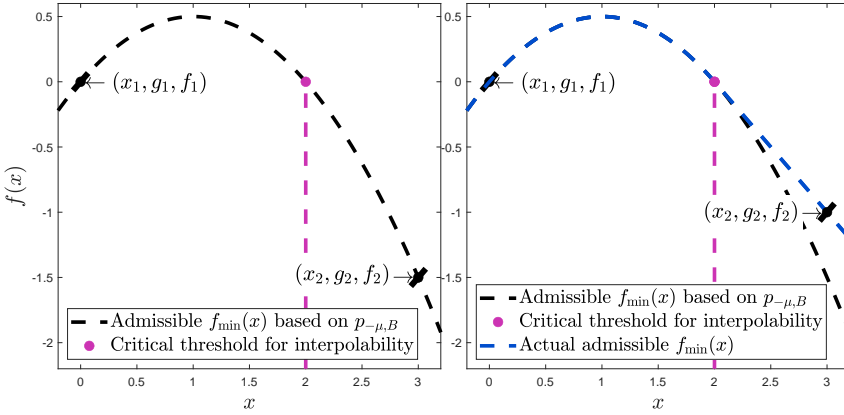


Fig. 4.1 Counterexample to the interpolability of $(p_{-\mu,B})$. On the left-hand side figure, the set $S = \{(x_i, f_i, g_i)\}_{i=1,2}$ satisfies $(p_{-\mu,B})$, but cannot be interpolated by a function in $\mathcal{F}_{p_{-\mu,B}}^m(\mathbb{R}^d)$. Indeed, the norm of the gradient associated to the only weakly convex function interpolating S , $f_{\min}(x) = -\frac{x^2}{2} + x$, grows larger than allowed from the critical threshold $x = 2$. On the right-hand side figure, the minimal function value allowed for f_2 ensures a bounded subgradient associated to the minimal weakly function interpolating S .

To strengthen $(p_{-\mu,B})$, consider $(\tilde{p}_{-\mu,B})$.

$$\tilde{p}_{-\mu,B}^{x,y} \leq 0 \Leftrightarrow \begin{cases} f(y) \geq f(x) + \langle g_x, y - x \rangle - \frac{\mu}{2} \|x - y\|^2 + \frac{\mu}{2} \|y - C_{xy}\|^2 \\ \|g_x\| \leq B, \\ f(x) \geq f_\star \end{cases} \quad (\tilde{p}_{-\mu,B})$$

where C_{xy} is the projection of y onto $\mathcal{B}_{\frac{g_x}{\mu} + x, \frac{B}{\mu}} := \{z \in \mathbb{R}^d : \|\mu(z - x) + g_x\| \leq B\}$, i.e.,

$$\begin{aligned} C_{xy} &= \operatorname{argmin}_{z \in \mathcal{B}_{\frac{g_x}{\mu} + x, \frac{B}{\mu}}} \|z - y\| & (4.12) \\ &= \begin{cases} \frac{g_x}{\mu} + x + \frac{B}{\mu} \frac{\mu(y-x) - g_x}{\|\mu(y-x) - g_x\|} & \text{if } \|\mu(y-x) - g_x\| \geq B, \\ y & \text{else.} \end{cases} \end{aligned}$$

As illustrated in Figure 4.1, $(\tilde{p}_{-\mu, B})$ strengthens $(p_{-\mu, B})$ by imposing the minimal function value allowed for f_j with respect to some (x_i, f_i, g_i) to always be compatible with a bounded subgradient. Observe that C_{xy} is the critical threshold for interpolability on Figure 4.1. Whenever two points are sufficiently close, e.g., any $x_2 \in (0, 2]$ in Figure 4.1, then $C_{xy} = y$ and $(\tilde{p}_{-\mu, B})$ reduces to $(p_{-\mu, B})$. However, once this critical threshold C_{xy} is reached, $(\tilde{p}_{-\mu, B})$ imposes a larger minimal function value.

Figure 4.2 illustrates the difference between the admissible regions allowed by $(p_{-\mu, B})$ and $(\tilde{p}_{-\mu, B})$.

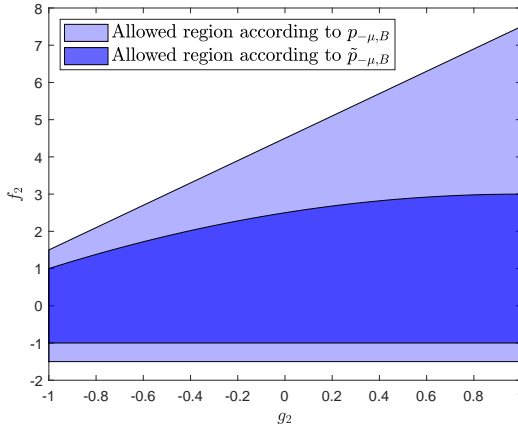


Fig. 4.2 Given $\mu = B = 1$, $(x_1, f_1, g_1) = (0, 0, 1)$, and $x_2 = 3$, allowed region for f_2 as a function of g_2 , according to $(p_{-\mu, B})$ and $(\tilde{p}_{-\mu, B})$. Clearly, $(p_{-\mu, B})$ is weaker than $(\tilde{p}_{-\mu, B})$ since satisfied by triplets that are not $\mathcal{F}_{p_{-\mu, B}}^m(\mathbb{R}^d)$ -interpolable.

4.5.3 Interpolation conditions for $\mathcal{F}_{p-\mu,B}^m(\mathbb{R}^d)$

This section formally proves that $(\tilde{p}_{-\mu,B})$, guessed from a graphical intuition, is indeed an interpolation constraint for $\mathcal{F}_{p-\mu,B}^m(\mathbb{R}^d)$.

Theorem 4.2 (Interpolation condition for $\mathcal{F}_{p-\mu,B}^m(\mathbb{R}^d)$)

The constraint $(\tilde{p}_{-\mu,B})$ is an interpolation constraint for $\mathcal{F}_{\tilde{p}_{-\mu,B}}^m(\mathbb{R}^d) = \mathcal{F}_{p-\mu,B}^m(\mathbb{R}^d)$.

We proceed to the proof of Theorem 4.2 by (i) proving that $(\tilde{p}_{-\mu,B})$ is an interpolation constraint, using pointwise extensibility, and (ii) showing that it characterizes the class of weakly convex functions with bounded subgradients when imposed everywhere.

Proposition 4.4 (Interpolability of $(\tilde{p}_{-\mu,B})$). *The constraint $(\tilde{p}_{-\mu,B})$ is an interpolation constraint for $\mathcal{F}_{\tilde{p}_{-\mu,B}}^m(\mathbb{R}^d)$.*

Proof. Consider $S = \{(x_i, f_i, g_i)\}_{i \in [N]}$ satisfying p , and any $z \in \mathbb{R}^d$. We seek for some $f_z \in \mathbb{R}$, $g_z \in \mathbb{R}^d$ extending S to z , that is

$$f_z \geq f_i + \langle g_i, z - x_i \rangle - \frac{\mu}{2} \|z - x_i\|^2 + \frac{\mu}{2} \|z - C_{i+}\|^2, \quad \forall i \in [N] \quad (4.13)$$

$$f_j \geq f_z + \langle g_z, x_j - z \rangle - \frac{\mu}{2} \|z - x_j\|^2 + \frac{\mu}{2} \|x_j - C_{+j}\|^2 := f_j^-, \quad \forall j \in [N] \quad (4.14)$$

$$\|g_z\| \leq B, \quad (4.15)$$

$$f_z \geq f_\star. \quad (4.16)$$

where C_{i+} denotes the projection of z onto $\mathcal{B}_{\frac{\mu}{2} + x_i, \frac{B}{\mu}}$ and C_{+j} denotes the projection of x_j onto $\mathcal{B}_{\frac{\mu}{2} + z, \frac{B}{\mu}}$. Let

$$i^* = \operatorname{argmax}_{i \in [N]} f_i + \langle g_i, z - x_i \rangle - \frac{\mu}{2} \|z - x_i\|^2 + \frac{\mu}{2} \|z - C_{i+}\|^2,$$

$$f_z = \max \left(f_{i^*} + \langle g_{i^*}, z - x_{i^*} \rangle - \frac{\mu}{2} \|z - x_{i^*}\|^2 + \frac{\mu}{2} \|z - C_{i^*+}\|^2, f_\star \right)$$

$$g_z = g_{i^*} + \mu(x_{i^*} - C_{i^*+}) \text{ if } f_z \neq f_\star, g_z = 0 \text{ otherwise.}$$

Then, f_z satisfies (4.13), (4.16), and, by definition of C_{i^*+} , g_z satisfies (4.15).

It remains to show (4.14) is satisfied, i.e., $\forall j \in [N], f_j \geq f_j^-$. Suppose first $f_z = f_*$. Then,

$$f_j^- = f_* - \frac{\mu}{2} \|z - x_j\|^2 + \frac{\mu}{2} \|x_j - C_{+j}\|^2 \leq f_* \leq f_j,$$

where the first inequality follows from $g_z = 0$ and the definition of C_{+j} . Suppose now $f_z \neq f_*$. Given the expression of f_z and g_z , f_j^- becomes

$$\begin{aligned} f_j^- &= f_{i^*} + \langle g_{i^*}, x_j - x_{i^*} \rangle + \mu \langle x_{i^*} - C_{i^*+}, x_j - z \rangle \\ &\quad - \frac{\mu}{2} \|z - x_j\|^2 + \frac{\mu}{2} \|x_j - C_{+j}\|^2 - \frac{\mu}{2} \|z - x_{i^*}^*\|^2 + \frac{\mu}{2} \|z - C_{i^*+}\|^2 \\ &= f_{i^*} + \langle g_{i^*}, x_j - x_{i^*} \rangle - \frac{\mu}{2} \|x_{i^*} - x_j\|^2 + \frac{\mu}{2} \|x_j - C_{+j}\|^2 \\ &\quad + \frac{\mu}{2} \|z - C_{i^*+}\|^2 + \mu \langle z - C_{i^*+}, x_j - z \rangle. \end{aligned}$$

By definition of g_z , $C_{i^*j} + z - C_{i^*+} \in \mathcal{B}_{\frac{g}{\mu} + x, \frac{B}{\mu}}$, implying

$$\begin{aligned} \|x_j - C_{+j}\|^2 &\leq \|x_j - C_{i^*j} - x + C_{i^*+}\|^2 \\ &= \|x_j - C_{i^*j}\|^2 + \|z - C_{i^*+}\|^2 \\ &\quad - 2 \langle x_j - C_{i^*j}, z - C_{i^*+} \rangle \\ \Leftrightarrow \|z - C_{i^*+}\|^2 + \|x_j - C_{+j}\|^2 &\leq \|x_j - C_{i^*j}\|^2 + 2 \|z - C_{i^*+}\|^2 \\ &\quad - 2 \langle x_j - C_{i^*j}, x - C_{i^*+} \rangle \\ &= \|z_j - C_{i^*j}\|^2 \\ &\quad - 2 \langle x_j - C_{i^*j} + C_{i^*+} - z, z - C_{i^*+} \rangle. \end{aligned}$$

Hence,

$$\begin{aligned} f_j^- &\leq f_{i^*} + \langle g_{i^*}, x_j - x_{i^*} \rangle - \frac{\mu}{2} \|x_{i^*} - x_j\|^2 + \frac{\mu}{2} \|x_j - C_{i^*j}\|^2 \\ &\quad + \mu \langle C_{i^*j} - C_{i^*+}, z - C_{i^*+} \rangle \tag{4.17} \\ &\leq f_{i^*} + \langle g_{i^*}, x_j - x_{i^*} \rangle - \frac{\mu}{2} \|x_{i^*} - x_j\|^2 + \frac{\mu}{2} \|x_j - C_{i^*j}\|^2, \end{aligned}$$

where the last inequality follows from (4.18). Finally, observe that $\forall j \in [N]$,

$$f_j \underset{p_{ii^*} \geq 0}{\geq} f_{i^*} + \langle g_{i^*}, x_j - x_{i^*} \rangle - \frac{\mu}{2} \|x_{i^*} - x_j\|^2 + \frac{\mu}{2} \|x_j - C_{i^*j}\|^2 \underset{(4.17)}{\geq} f_j^-,$$

so that (4.14) is satisfied. Hence, $(\tilde{p}_{-\mu,B})$ is pointwise extensible. In addition, $(\tilde{p}_{-\mu,B})$ is continuous and differentially consistent since $-\frac{\mu}{2}x^2 = \omega(x)$, and $\|g\| \leq B$, i.e., g is regularly extensible since it satisfies (4.7), see Lemma 4.1. By Theorem 4.1, it is an interpolation constraint for $\mathcal{F}_{\tilde{p}_{-\mu,B}}^m(\mathbb{R}^d)$. \square

Proposition 4.5 (Equivalence of $\mathcal{F}_{p_{-\mu,B}}^m(\mathbb{R}^d)$ and $\mathcal{F}_{\tilde{p}_{-\mu,B}}^m(\mathbb{R}^d)$). *Let $(p_{-\mu,B})$ and $(\tilde{p}_{-\mu,B})$. Then $\mathcal{F}_{p_{-\mu,B}}^m(\mathbb{R}^d) = \mathcal{F}_{\tilde{p}_{-\mu,B}}^m(\mathbb{R}^d)$.*

The proof of Proposition 4.5 resorts to the following Lemma.

Lemma 4.3. *Let y_0 be the projection onto $\mathcal{B}_{c,R}$ of $y \in \mathbb{R}^d \notin \mathcal{B}_{c,R}$, for some $c \in \mathbb{R}^d$ and $R \in \mathbb{R}_+$. Let $z = \alpha y + (1 - \alpha)y_0$ for some $\alpha \in [0, 1]$. Then, y_0 is also the projection onto $\mathcal{B}_{c,R}$ of z . In addition, $\forall w \in \mathbb{R}^d$ such that $\|w\| \leq R$,*

$$\langle w, z - y \rangle \leq \langle c - y_0, z - y \rangle.$$

Proof. Given any closed set C and any $y \in \mathbb{R}^d$, y_0 is the projection of y onto C if and only if

$$\forall x \in \mathbb{R}^d, \langle x - y_0, y - y_0 \rangle \leq 0 \quad (4.18)$$

see, e.g., [Ber09, Proposition 1.1.9]. Therefore, $\forall x \in \mathcal{B}_{c,R}$, $\langle x - y_0, z - y_0 \rangle = \alpha \langle x - y_0, y - y_0 \rangle \stackrel{(4.18)}{\leq} 0$. Hence, y_0 is the projection of z onto $\mathcal{B}_{c,R}$. It also holds that for any $\|w\| \leq R$, $\langle w, z - y \rangle \leq R\|z - y\|$, and $c - y_0$ reaches that since of norm R , and of same direction as $z - y$. Hence, $c - y_0 = \frac{R(z-y)}{\|z-y\|}$. \square

We now proceed to the proof of Proposition 4.5.

Proof. The case $\mu = 0$ and the inclusion $\mathcal{F}_{\tilde{p}_{-\mu,B}}^m(\mathbb{R}^d) \subseteq \mathcal{F}_{p_{-\mu,B}}^m(\mathbb{R}^d)$ follow from the definition, since the inequalities only differ by a positive term $\mu\| \cdot \|^2$. Hence, we consider $\mu > 0$ and establish $\mathcal{F}_{p_{-\mu,B}}^m(\mathbb{R}^d) \subseteq \mathcal{F}_{\tilde{p}_{-\mu,B}}^m(\mathbb{R}^d)$.

By contradiction, suppose there exists a mapping (f, g) satisfying everywhere $(p_{-\mu,B})$ but not $(\tilde{p}_{-\mu,B})$. Then, at some $x, y \in \mathbb{R}^d$, given $f_x = f(x)$, $f_y = f(y)$, and for some $g_x \in g(x)$, $g_y \in g(y)$, $(p_{-\mu,B})$ is satisfied while $(\tilde{p}_{-\mu,B})$ is violated, i.e.,

$$p_{-\mu,B}^{xy} \leq 0, p_{-\mu,B}^{yx} \leq 0, \tilde{p}_{-\mu,B}^{xy} > 0, \text{ w.l.o.g. in the choice of } x \text{ and } y.$$

Since $p_{-\mu,B}^{xy}$ and $\tilde{p}_{-\mu,B}^{xy}$ only differ by $\mu\|C_{xy} - y\|^2$, this implies $C_{xy} \neq y$, where C_{xy} is the projection of y onto $\mathcal{B}_{\frac{g_x}{\mu} + x, \frac{g}{\mu}}$, and the existence of some

$\beta_0 \in [0, 1)$ such that

$$f_y = f_x + \langle g_x, y - x \rangle - \frac{\mu}{2} \|x - y\|^2 + \frac{\beta_0 \mu}{2} \|C_{xy} - y\|^2, \quad (4.19)$$

since f_y is larger than the right-hand side with $\beta_0 = 0$ due to $p_{-\mu, B}^{xy} \leq 0$, and smaller than the right-hand side with $\beta_0 = 1$ due to $\tilde{p}_{-\mu, B}^{xy} > 0$. Consider the decreasing sequence (β_k) , initialized in (4.19) and defined by

$$\beta_{k+1} = \begin{cases} 2 - \frac{1}{\beta_k} & \beta_k > \frac{1}{2} \\ 0 & \beta_k \leq \frac{1}{2}, \end{cases} \quad (4.20)$$

and the associated sequence (z_k) , initialized in $z_0 = y$ and defined by

$$z_{k+1} = \begin{cases} \beta_k z_k + (1 - \beta_k) C_{xy}, & \beta_k \neq 0, \\ \frac{1}{2}(y + C_{xy}), & \beta_k = 0. \end{cases}$$

Let $k = 1$. Since (f, g) satisfy $(p_{-\mu, B})$ everywhere, letting $f_{z_1} = f(z_1)$, and for all $g_{z_1} \in g(z_1)$, it holds,

$$\begin{aligned} f_{z_1} &\geq f_x + \langle g_x, z_1 - x \rangle - \frac{\mu}{2} \|z_1 - x\|^2 := f_{z_1}^-, \\ f_{z_1} &\leq f_y + \langle g_{z_1}, z_1 - y \rangle + \frac{\mu}{2} \|z_1 - y\|^2 := f_{z_1}^+, \\ \|g_{z_1}\| &\leq B. \end{aligned} \quad (4.21)$$

Since C_{xy} is the projection of y onto $\mathcal{B}_{\frac{\mu}{2} + x, \frac{B}{\mu}}$ and z_1 is a convex combination of C_{xy} and y , by Lemma 4.3, it holds

$$\langle g_{z_1}, z_1 - y \rangle \leq \langle g_x + \mu(x - C_{xy}), z_1 - y \rangle. \quad (4.22)$$

In addition, by definition, $z_1 - y = -K(z_1 - C_{xy})$, where $K = \frac{(1 - \beta_0)}{\beta_0}$ if $\beta_0 \neq 0$, $K = 1$ otherwise. We have

$$\begin{aligned} \langle x - C_{xy}, z_1 - y \rangle &= \langle z_1 - C_{xy}, z_1 - y \rangle + \langle x - z_1, z_1 - y \rangle \\ &= -K \|z_1 - C_{xy}\|^2 + \frac{\|x - y\|^2}{2} - \frac{\|z_1 - x\|^2}{2} - \frac{\|z_1 - y\|^2}{2}. \end{aligned} \quad (4.23)$$

Consequently,

$$\begin{aligned}
 f_{z_1}^+ &\stackrel{(4.22)}{\leq} f_y + \langle g_x, z_1 - y \rangle + \mu \langle x - C_{xy}, z_1 - y \rangle + \frac{\mu \|z_1 - y\|^2}{2} \\
 &\stackrel{(4.23)}{=} f_y + \langle g_x, z_1 - y \rangle - K \|z_1 - C_{xy}\|^2 + \frac{\mu \|x - y\|^2}{2} - \frac{\mu \|z_1 - x\|^2}{2} \\
 &\stackrel{(4.19)}{=} f_x + \langle g_x, z_1 - x \rangle - \frac{\mu}{2} \|z_1 - x\|^2 + \frac{\mu \beta_0}{2} \|C_{xy} - y\|^2 - K \|C_{xy} - z_1\|^2
 \end{aligned}$$

If $\beta_0 = 0$, this implies $f_{z_1}^+ = f_{z_1}^- - \|C_{xy} - z_1\|^2 < f_{z_1}^-$, hence the assumption " (f, g) satisfy $(p_{-\mu, B})$ everywhere" cannot hold at z_1 . Else,

$$f_{z_1}^+ \stackrel{C_{xy}-y=}{=} f_{z_1}^- + \frac{\mu}{2} \frac{2\beta_0 - 1}{\beta_0} \|C_{xy} - z_1\|^2 = f_{z_1}^- + \frac{\mu}{2} \frac{2\beta_0 - 1}{\beta_0} \|C_{xz_1} - z_1\|^2,$$

$\frac{1}{\beta_0} (C_{xy} - z_1)$

where the last equality follows from the first assertion in Lemma 4.3. By definition of β_1 , $f_{z_1}^-$ and $f_{z_1}^+$, it then holds

$$f_{z_1} = f_x + \langle g_x, z_1 - x \rangle - \frac{\mu}{2} \|z_1 - x\|^2 + \frac{\beta_1 \mu}{2} \|C_{xz_1} - z_1\|^2, \quad (4.24)$$

i.e., (4.19) with z_1 and β_1 replacing z_0 and β_0 . Starting from (4.24) and repeating the argument iteratively, it holds, $\forall k \geq 1$,

$$f_{z_k} = f_x + \langle g_x, z_k - x \rangle - \frac{\mu}{2} \|z_k - x\|^2 + \frac{\beta_k \mu}{2} \|C_{xz_k} - z_k\|^2, \quad (4.25)$$

where β_k is defined as in (4.20). The sequence (β_k) can be equivalently defined as

$$\beta_k = \begin{cases} \frac{(k+1)\beta_0 - k}{k\beta_0 - (k-1)} & k < \frac{2\beta_0 - 1}{1 - \beta_0} \\ 0 & k \geq \frac{2\beta_0 - 1}{1 - \beta_0}. \end{cases} \quad (4.26)$$

Hence, at any z_{k+1} where $k \geq \frac{2\beta_0 - 1}{1 - \beta_0}$, following the same argument as when $\beta_0 = 0$, replacing β_0 by $\beta_k = 0$ and z_1 by z_{k+1} , the assumption " (f, g) satisfy $(p_{-\mu, B})$ everywhere" cannot hold, which completes the proof. \square

Combining Propositions 4.4 and 4.5 proves Theorem 4.2.

4.5.4 A Gram-representable formulation of $(\tilde{p}_{-\mu,B})$

Our goal is to combine the PEP framework, introduced in Chapter 3, with $(\tilde{p}_{-\mu,B})$, the exact description of $\mathcal{F}_{\tilde{p}_{-\mu,B}}^m(\mathbb{R}^d)$, to tightly analyze methods on this function class. As discussed in Chapter 3, PEP formulations can be solved efficiently, provided all constraints characterizing the problem class are Gram-representable (Definition 3.1), i.e., linear in function values and scalar products of iterates and (sub)gradients. As such, given the expression of C_{ij} given in (4.12), $(\tilde{p}_{-\mu,B})$ is not Gram-representable. However, we show that it can be equivalently expressed in a Gram-representable way.

Lemma 4.4 (Gram-representable formulation of $(\tilde{p}_{-\mu,B})$). *Consider $(\tilde{p}_{-\mu,B})$. Then, for any $x_i, x_j \in \mathbb{R}^d$,*

$$\tilde{p}_{-\mu,B}^{i,j} \leq 0 \Leftrightarrow \begin{cases} \exists C_{ij} \in \mathbb{R}^d : \\ f_j \geq f_i + \langle g_i, x_j - x_i \rangle - \frac{\mu}{2} \|x_i - x_j\|^2 + \frac{\mu}{2} \|x_j - C_{ij}\|^2, \\ \|g_i + \mu(x_i - C_{ij})\|^2 \leq B^2, \\ \|g_i\|^2 \leq B^2, \\ f_z \geq f_*. \end{cases} \quad (4.27)$$

Proof. Denote by C_{ij}^* the projection of x_j onto $\mathcal{B}_{\frac{g_i}{\mu} + x_i, \frac{B}{\mu}}$. Then,

$$\exists C_{ij} \in \mathcal{B}_{\frac{g_i}{\mu} + x_i, \frac{B}{\mu}} : f_j \geq f_i + \langle g_i, x_j - x_i \rangle - \frac{\mu}{2} \|x_i - x_j\|^2 + \frac{\mu}{2} \|x_j - C_{ij}\|^2 \quad (4.28)$$

$$\Leftrightarrow f_j \geq f_i + \langle g_i, x_j - x_i \rangle - \frac{\mu}{2} \|x_i - x_j\|^2 + \frac{\mu}{2} \|x_j - C_{ij}^*\|^2. \quad (4.29)$$

Indeed, (4.29) clearly implies (4.28). In addition, if (4.28) is satisfied, since $\|x_j - C_{ij}\| \geq \|x_j - C_{ij}^*\|$, (4.29) is satisfied. \square

4.5.5 Performance analysis of the subgradient method on $\mathcal{F}_{p_{-\mu,B}}^m(\mathbb{R}^d)$.

Combining the PEP framework and the interpolation condition $(\tilde{p}_{-\mu,B})$, under its reformulation (4.27), allows tightly analyzing the subgradient method on $\mathcal{F}_{p_{-\mu,B}}^m(\mathbb{R}^d)$. The subgradient method's iteration is given by

$$x_{k+1} = x_k - \alpha g_k, \quad g_k \in \partial f(x_k), \quad (4.30)$$

In non-convex non-smooth optimization, usual performance criteria such as $f(x_N) - f_*$, or $\text{dist}(0, \partial f(x_N))$ may not tend to 0, and may vary discontinuously [DD19, Nur73]. Hence, classical performance criteria in this setting rely on a continuously differentiable approximation of f , that is, its *Moreau envelope* [Mor65], given by

$$f_\rho(x) := \min_{y \in \mathbb{R}^d} f(y) + \frac{\|x - y\|^2}{2\rho},$$

where $\rho \in (0, \frac{1}{\mu}]$ [DDMP18]. Specifically, given $(x_k)_{k=0,1,\dots,N}$ the iterates of (4.30), we consider as performance measure

$$\mathcal{P} = \frac{1}{N+1} \sum_{i=0}^N \|\nabla f_\rho(x_i)\|^2 = \frac{1}{N+1} \sum_{i=0}^N \frac{\|x_i - \mathbf{prox}_{\rho f}(x_i)\|^2}{\rho^2},$$

where $\mathbf{prox}_{\rho f}(x) := \underset{y \in \mathbb{R}^d}{\text{argmin}} f(y) + \frac{\|x - y\|^2}{2\rho} \quad \forall x \in \mathbb{R}^d,$

and as initial condition,

$$f_\rho(x_0) - f^* \leq R^2 \Leftrightarrow f_0^y - f^* + \frac{\|x_0 - y_0\|^2}{2\rho} \leq R^2.$$

This performance criterion is a measure of near-stationarity for f , since $\text{dist}(0, \partial f(x)) \leq \|\nabla f_\rho(x)\|$ [DD19, Section 2.2]. The PEP yielding the exact convergence rate on this problem is

$$\max_{\substack{S = \{(x_i, f_i, g_i)\}_{i=0,1,\dots,N,*} \\ S_y = \{(y_i, f_i^y, g_i^y)\}_{i=0,1,\dots,N}}} \frac{1}{N+1} \sum_{i=0}^N \frac{\|x_i - y_i\|^2}{\rho^2}$$

Subgradient method:

$$\text{s.t. } x_{i+1} = x_i - h g_i, \quad i = 0, 1, \dots, N-1$$

$$y_i = \mathbf{prox}_{\rho f}(x_i):$$

$$y_i = x_i - \rho g_i^y, \quad i = 0, 1, \dots, N$$

$S \cup S_y$ is $\mathcal{F}_{\tilde{p}_{-\mu, B}}^m(\mathbb{R}^d)$ -interpolable:

$$\tilde{p}_{-\mu, B}^{ij} \leq 0, \quad \forall i, j \in S \cup S_y$$

Initial condition:

$$f_0^y - f^* + \frac{\|x_0 - y_0\|^2}{2\rho} \leq R^2$$

Optimality:

$$g^* = 0.$$

(4.31)

Relying on (4.27) to enforce $\mathcal{F}_{\tilde{p}_{-\mu, B}}^m(\mathbb{R}^d)$ -interpolability of $S \cup S_y$, all con-

straints and objective function in (4.31) are Gram-representable, hence this PEP can be efficiently solved.

Setting $\rho = \frac{1}{2\mu}$ as in [DDMP18, DGVPR24], we compare the rate resulting from (4.27) with two existing convergence rates from the literature. First, [DD19] analyzed (4.30) on $\mathcal{F}_{\bar{p}-\mu, B}^m(\mathbb{R}^d)$ in a classical way, using $(p_{-\mu, B})$, and showed a convergence rate [DD19, Equation (3.4)] of

$$\mathcal{P} \leq 2 \frac{R^2 + \mu B^2 (N+1) \alpha^2}{(N+1) \alpha}, \quad (4.32)$$

whose stepsize is optimized in [DD19, Theorem 3.1]

$$\alpha = \frac{R}{B \sqrt{\mu} \sqrt{N+1}}, \quad (4.33)$$

which we refer to as classical stepsize. Then, [DGVPR24] analyzed the same setting, in a computer-aided way. The authors relied on a Lyapunov function approach, analyzing a single iteration of (4.30) with respect to a specific Lyapunov structure, and using $(p_{-\mu, B})$. They obtained a rate [DGVPR24, Equation (45)] of

$$\mathcal{P} \leq \frac{B^2 \left(-1 + (1 - 2\alpha)^{N+1} + 2\alpha(N+1) + \frac{2R^2\mu^2}{\alpha B^2} (1 - (1 - 2\alpha)^{N+1}) \right)}{\mu^2 (N+1)}, \quad (4.34)$$

whose stepsize is optimized in [DGVPR24, Corollary 1]

$$\alpha = \frac{\sqrt{4 \left(\frac{R\mu}{B} \right)^2 (N+1) + 1}}{2(N+1)}, \quad (4.35)$$

which we refer to as efficient stepsize.

Remark 4.2. By boundedness of the subgradient and definition of \mathcal{P} , the performance measure is of at most B^2 . Hence, bounds (4.32) and (4.34) become non-trivial from $N \geq \left(\frac{4R}{\mu B}\right)^2 - 1$ and $N \geq \left(\frac{2R}{\mu B}\right)^4 - 1$, respectively.

Figure 4.3 compares (4.32), (4.34), a PEP-based bound using the non-tight constraint $(p_{-\mu, B})$, and (4.31), i.e., the tightest bound possible. Both PEP-based bounds outperform the other ones, and are always non-trivial. On this example, the effect of optimally combining all constraints (computing PEP-based bounds) is more impactful than that of using interpola-

tion constraints. Interestingly, the efficient stepsize (4.35) as computed in [DGVPR24] turns out to yield a worst convergence rate than the stepsize (4.33) derived in [DD19], highlighting the importance of tuning methods on the basis of tight analyses.




Fig. 4.3 Comparison of worst-case performance guarantees for (4.30) on $\mathcal{F}_{\tilde{p}_{-\mu,B}}^m$, with parameters $f_\star = 0$, $R^2 = \frac{1}{8}$ and $B, \mu = 1$. The top figure displays the evolution, with the number of iterations, of [DD19, Equation (3.4)] with classical stepsize (4.33) (pink); of [DGVPR24, Equation (45)] with efficient stepsize (4.35) (black), and of the PEP-based exact bounds relying on $(p_{-\mu,B})$ and $(\tilde{p}_{-\mu,B})$ (blue), with stepsize (4.33). The bottom figure displays the evolution of these bounds with the stepsize α , after 5 iterations.

4.6 Exhaustive analysis of a class of constraints

Motivated by the difficulty to obtain equivalent characterizations of given a function class, and by the observation that function classes considered as interesting in optimization do not always have a simple tight description, see Section 4.4, we illustrate the constraint-based approach to interpolation presented in Section 4.2 on a simple class of constraints. We consider constraints consisting of a single inequality, linear in function values and scalar product of points and subgradients, and such that $p^{xx} = 0$, i.e., any single data point satisfies p . Hence, we analyze constraints

$$p^{ij} \leq 0 \Leftrightarrow f_i \geq f_j + B\|g_i\|^2 + C\|g_j\|^2 + D\langle g_j, g_i \rangle + E\|x_i\|^2 + F\|x_j\|^2 \\ + G\langle x_i, x_j \rangle + H\langle g_i, x_i \rangle + I\langle g_i, x_j \rangle + J\langle g_j, x_i \rangle + K\langle g_j, x_j \rangle, \quad (4.36)$$

where $D = -(B + C)$, $G = -(E + F)$ and $K = -(H + I + J)$ to satisfy $p^{xx} = 0$. Such constraints are Gram representable (Definition 3.1), and tightly describe several first-order function classes, e.g., μ -strongly convex L -smooth functions (Proposition 2.1).

We show that the only differentially consistent interpolation constraints of the form (4.36) are (2.15), characterizing $\mathcal{F}_{\mu, L}$, and a constraint describing the class of quadratic functions $f(x) = \frac{\mu}{2}\|x\|^2 + bx + c$ for some $b \in \mathbb{R}^d$, $c \in \mathbb{R}$, i.e., the limit case of (2.15). Equivalently, these are the only function classes described by (4.36) that admit a simple tight characterization.

Theorem 4.3 (Exhaustive analysis of a class of constraints)

Let p be of the form (4.36). Then p is (i) differentially consistent, and (ii) an interpolation constraint if and only if $p^{ij} \leq 0$ can be expressed as in (2.15), where $-\infty < \mu \leq L \leq +\infty$ and $L \geq 0$, or as

$$p^{ij} \leq 0 \Leftrightarrow f_i \geq f_j + \frac{1}{2}\langle g_i + g_j, x_i - x_j \rangle + M\|g_i - g_j - \mu(x_i - x_j)\|^2, \quad (4.37)$$

where $M > 0$. In the first case, $\mathcal{F}_p^m(\mathbb{R}^d)$ is the class of μ - (strongly/weakly) convex L -smooth functions, and in the second case, $\mathcal{F}_p^m(\mathbb{R}^d)$ is the class of quadratic functions $f(x) = \frac{\mu}{2}\|x\|^2 + bx + c$ for some $b \in \mathbb{R}^d$, $c \in \mathbb{R}$.

Remark 4.3. The second function class is in itself trivial, since it can be exactly characterized by equalities $g_i = \mu x_i + b$, $f_i = \frac{\mu}{2} \|x_i\|^2 + bx + c$. However, we present its characterization via a constraint of the form (4.36) (i) as a part of the exhaustive analysis we conduct, and (ii) to characterize the whole spectrum $-\infty < \mu \leq L \leq +\infty$ by a set of similar inequalities, without having to distinguish $\mu = L$ as a special case.

To prove Theorem 4.3, we first show that (2.15) and (4.37) are indeed differentially consistent interpolation constraints. The case (2.15) follows from Proposition 2.1 and [RGP22, Theorem 3.1] for the case $\mu < 0$, and Proposition 4.6 tackles (4.37).

Proposition 4.6. *Consider a constraint p defined by (4.37). Then p is a differentially consistent interpolation constraint, and the class $\mathcal{F}_p^m(\mathbb{R}^d)$ it defines is the class of quadratic functions of the form $f(x) = \frac{\mu}{2} \|x\|^2 + bx + c$, for some $b \in \mathbb{R}^d$, $c \in \mathbb{R}$.*

Proof. First, note that $p^{ij} \leq 0$ implies

$$\begin{aligned} f_i &\geq f_j + \langle g_j, x_i - x_j \rangle + M \|g_i - g_j + \frac{1 - 4M\mu}{4M} (x_i - x_j)\|^2 \\ &\quad + (M\mu^2 - \frac{(1 - 4M\mu)^2}{16M}) \|x_i - x_j\|^2 \\ &\geq f_j + \langle g_j, x_i - x_j \rangle + \omega(\|x_i - x_j\|). \end{aligned}$$

Hence, $p^{ij} \leq 0$ implies (4.4) and p is differentially consistent.

Consider now $S = \{(x_i, f_i, g_i)\}_{i \in [N]}$ satisfying p , and any $x \in \mathbb{R}^d$. For any $i, j \in [N]$, summing $p^{ij} \leq 0$ and $p^{ji} \leq 0$ implies

$$0 \geq \|(g_i - g_j) - \mu(x_i - x_j)\|^2 \Leftrightarrow g_i - g_j = \mu(x_i - x_j), \quad (4.38)$$

which in turn implies, by combining $p^{ij} \leq 0$ and $p^{ji} \leq 0$,

$$f_i = f_j + \frac{1}{2} \langle g_i + g_j, x_i - x_j \rangle. \quad (4.39)$$

We let $g = g_i + \mu(x - x_i)$ for an arbitrary $i \in [N]$, which implies by (4.38) $g = g_j + \mu(x - x_j)$, $\forall j \in [N]$. Then, $M \|g - g_i - \mu(x - x_i)\|^2 = 0$, and extending p to x amounts to obtaining a f such that $f = f_i + \frac{1}{2} \langle g + g_i, x - x_i \rangle$, $\forall i \in [N]$. Let

$$f = f_i + \frac{1}{2} \langle g + g_i, x - x_i \rangle$$

$$\begin{aligned}
 &\stackrel{(4.39)}{=} f_j + \frac{1}{2} \langle g + g_j, x - x_j \rangle + \frac{1}{2} (\langle g, x_j - x_i \rangle + \langle g_i, x - x_j \rangle + \langle g_j, x_i - x \rangle) \\
 &\stackrel{(4.38)}{=} f_j + \frac{1}{2} \langle g + g_j, x - x_j \rangle \\
 &\quad + \frac{1}{2} (\langle g, x_j - x_i \rangle + \langle g - \mu(x - x_i), x - x_j \rangle + \langle g - \mu(x - x_j), x_i - x \rangle) \\
 &= f_j + \frac{1}{2} \langle g + g_j, x - x_j \rangle \quad \forall i, j \in [N].
 \end{aligned}$$

This choice of f and g extends p to x , hence p is pointwise extensible. Since $\|g\| \leq \|g_i\| + \mu\|x - x_i\|$, $\forall i \in [N]$, and p is continuous and differentially consistent, p is also regularly pointwise extensible, since it satisfies (4.6), see Lemma 4.1. Hence, by Theorem 4.1, p is an interpolation constraint.

Finally, it holds that $\mathcal{F}_p^m(\mathbb{R}^d)$ is the class of quadratic functions $f = \frac{\mu}{2}\|x\|^2 + bx + c$. Indeed, consider any $f \in \mathcal{F}_p(\mathbb{R}^d)$. Then, by (4.38) and (4.39), it holds $\forall x, y \in \mathbb{R}^d$ and $g(x) = \nabla f(x)$,

$$\begin{aligned}
 g(x) &= g(y) + \mu(x - y) && \Leftrightarrow g(x) = \mu x + b, \\
 f(x) &= f(y) + \frac{\mu}{2}(\|x\|^2 - \|y\|^2) + b(x - y) && \Leftrightarrow f(x) = \frac{\mu}{2}\|x\|^2 + bx + c,
 \end{aligned}$$

for some $b \in \mathbb{R}^d$, $c \in \mathbb{R}$. Suppose now $f(x) = \frac{\mu}{2}\|x\|^2 + bx + c$. Then, $\forall x, y \in \mathbb{R}^d$, f satisfies p , hence $f \in \mathcal{F}_p^m(\mathbb{R}^d)$. \square

We now prove that no other constraint of the form (4.36) is a differentially consistent interpolation constraint. First, we analyze differential consistency of (4.36).

Proposition 4.7 (Differential consistency of (4.36)). *Consider a constraint p of the form (4.36). Then p is differentially consistent if and only if $p^{ij} \leq 0$ can be expressed as*

$$f_i \geq f_j + \langle \gamma g_i + (1 - \gamma)g_j, x_i - x_j \rangle + \alpha \|g_i - g_j\|^2 + \beta \|x_i - x_j\|^2, \quad (4.40)$$

where $\alpha \geq 0$, and $\gamma, \beta \in \mathbb{R}$.

Proof. First, note that given p of the form (4.40), $p^{ij} \leq 0$ implies

$$\begin{aligned}
 f_i &\geq f_j + \langle g_j, x_i - x_j \rangle + \alpha \|g_i - g_j + \frac{\gamma}{2\alpha}(x_i - x_j)\|^2 + (\beta - \frac{\gamma^2}{4\alpha}) \|x_i - x_j\|^2 \\
 &\geq f_j + \langle g_j, x_i - x_j \rangle + (\beta - \frac{\gamma^2}{4\alpha}) \|x_i - x_j\|^2
 \end{aligned}$$

$$= f_j + \langle g_j, x_i - x_j \rangle + \omega(\|x_i - x_j\|).$$

Hence, $p^{ij} \leq 0$ implies (4.4) and p is differentially consistent.

Consider now a constraint p of the form (4.36) but that cannot be written as in (4.40). Consider $f : \mathbb{R}^d \rightarrow \mathbb{R}$ and $g : \mathbb{R}^d \rightarrow \mathbb{R}^d$ satisfying p everywhere, and suppose that, at some $y \in \mathbb{R}^d$, f is twice differentiable. For any $\varepsilon \in \mathbb{R}^d$, let $x = y + \varepsilon$, and consider the Taylor development of order 1 of f and g at x :

$$f(x) = f_y + \langle \varepsilon, f'_y \rangle + \mathcal{O}(\varepsilon^2), \quad g(x) = g_y + g'_y \varepsilon + \mathcal{O}(\varepsilon^2),$$

where f_y, g_y, f'_y and g'_y denote, respectively, $f(y), g(y), \nabla f(y)$ and $\nabla^2 f(y)$. Satisfaction of p^{xy} then implies:

$$\begin{aligned} f_y + \langle \varepsilon, f'_y \rangle + \mathcal{O}(\varepsilon^2) &\geq f_y \\ &+ B(\|g_y\|^2 + 2\langle g_y, g'_y \varepsilon \rangle) - (B + C)(\|g_y\|^2 + \langle g'_y, g'_y \varepsilon \rangle) \\ &+ E(\|y\|^2 + 2\langle y, \varepsilon \rangle) + F\|y\|^2 - (E + F)(\|y\|^2 + \langle y, \varepsilon \rangle) \\ &+ H(\langle g_y, y \rangle + \langle g_y, \varepsilon \rangle + \langle g'_y \varepsilon, y \rangle) + I(\langle g_y, y \rangle + \langle g'_y \varepsilon, y \rangle) \\ &+ J(\langle g_y, y \rangle + \langle g_y, \varepsilon \rangle) - (I + J + K)\langle g_y, y \rangle + \mathcal{O}(\varepsilon^2) \\ &\Leftrightarrow \langle \varepsilon, f'_y \rangle \geq (B - C)\langle g_y, g'_y \varepsilon \rangle + (E - F)\langle y, \varepsilon \rangle \\ &+ (H + J)\langle g_y, \varepsilon \rangle + (H + I)\langle g'_y \varepsilon, y \rangle + \mathcal{O}(\varepsilon^2). \end{aligned}$$

Since this holds for any choice of ε , it implies for any $k \in [d]$:

$$f'_y{}^{(k)} = (H + J)g_y^{(k)} + (B - C)(g'_y g_y)^{(k)} + (E - F)y^{(k)} + (H + I)(g'_y y)^{(k)}.$$

Unless $B = C$, $E = F$, $H = -I$, and $H + J = 1$, i.e., the parameters corresponding to (4.40), $g_y \neq f'_y$. \square

Second, we show that if p is of the form (4.40), but cannot be expressed as in Theorem 4.3, then p is not an interpolation constraint.

Proposition 4.8 (Interpolability of (4.36)). *Consider a constraint p of the form (4.40). If p cannot be expressed as (2.15) or (4.37), then p is not pointwise extensible.*

Proof. Let $\alpha = 0$. Then, (4.40) can be written as (2.15) or (4.37) if and only if $\gamma = 0$ (corresponding to $L = \infty$) or $\gamma = 1$ (corresponding to $\mu = -\infty$). Table 4.2 provides counterexamples to the extensibility of (4.40) for any other γ .

| γ | Counterexample |
|-------------------------------|--|
| $\gamma < 0$ | $S = \left\{ \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}, \begin{pmatrix} -1 \\ \frac{\beta}{1-2\gamma} - 1 + \gamma \\ \frac{-2\beta}{1-2\gamma} \end{pmatrix}, \begin{pmatrix} -1 \\ \frac{\beta}{1-2\gamma} - 1 + \gamma \\ \frac{-2\beta}{1-2\gamma} \end{pmatrix} \right\}, z = -0.5$ |
| $\gamma > 1$ | $S = \left\{ \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}, \begin{pmatrix} -1 \\ \frac{\beta}{1-2\gamma} - 1 - \gamma \\ \frac{-2\beta}{1-2\gamma} + 2 \end{pmatrix} \right\}, z = -0.5$ |
| $0 < \gamma \leq \frac{1}{2}$ | $S = \left\{ \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}, \begin{pmatrix} -1 \\ \frac{\beta}{1-2\gamma} - 1 + \gamma \\ \frac{-2\beta}{1-2\gamma} \end{pmatrix}, \begin{pmatrix} -1.1 \\ \frac{\beta}{(1-2\gamma)} - 1 + \gamma \\ \frac{-2\beta}{1-2\gamma} + 1 \end{pmatrix} \right\}, z = -0.5$ |
| $\frac{1}{2} < \gamma < 1$ | $S = \left\{ \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}, \begin{pmatrix} -1 \\ \frac{\beta}{1-2\gamma} - 1 - \gamma \\ \frac{-2\beta}{1-2\gamma} + 2 \end{pmatrix}, \begin{pmatrix} -1.1 \\ \frac{\beta}{(1-2\gamma)} - 1 - \gamma \\ \frac{-2\beta}{1-2\gamma} + 1 \end{pmatrix} \right\}, z = -0.5$ |

Table 4.2 Counterexamples to the pointwise extensibility of (4.40) when $\alpha = 0$, $\gamma \neq 0, 1$. In the counterexamples, $S = \{(x_i, f_i, g_i)\}$, and z is the new point at which extending p is infeasible. One can verify these counterexamples using the code referenced in Section 4.1, which was used to gain intuition about them.

Let $\alpha > 0$, and consider $\Delta = \frac{\beta}{\alpha} + \frac{(1-\gamma)\gamma}{4\alpha^2}$. Then, (4.40) can be written as (2.15) or (4.37) if and only if $\Delta = 0$, leading to (2.15) or $\Delta = \frac{1}{16\alpha^2}$, leading to (4.37). Table 4.3 provides counterexamples to the interpolability of any other choice of Δ .

These counterexamples are one-dimensional, but they can straightforwardly be extended to any dimension by considering vectors with one non-zero entry.

| β | Counterexample |
|------------------------------------|---|
| $\beta < 0$ | $S = \left\{ \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 \\ -\beta + (1 - \gamma) \frac{1-2\gamma}{4\alpha} - \left(\frac{1-2\gamma}{4\alpha}\right)^2 \\ \frac{1-2\gamma}{4\alpha} \end{pmatrix} \right\},$ $z = 0.5$ |
| $\beta > \frac{1}{16\alpha^2}$ | $S = \left\{ \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 \\ -\beta + (1 - \gamma) \frac{1-2\gamma}{4\alpha} - \left(\frac{1-2\gamma}{4\alpha}\right)^2 \\ \frac{1-2\gamma}{4\alpha} \end{pmatrix} \right\},$ $z = \frac{1}{16\alpha^2\beta - 1} + 1$ |
| $0 < \beta < \frac{1}{16\alpha^2}$ | $S = \left\{ \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 4\alpha \\ 1 - 4\gamma + 2\gamma^2 + (1 - \gamma)\sqrt{K} \\ 1 - 2\gamma + \sqrt{K} \end{pmatrix} \right\},$ $z = 2\alpha(1 + K^{-1}), K = (1 - 2\gamma)^2 - 16\alpha\beta$ |

Table 4.3 Numerical counterexamples to the interpolability of (4.40) when $\alpha > 0$, $\Delta \neq \{0, \frac{1}{16\alpha^2}\}$. In the counterexamples, $S = \{(x_i, f_i, g_i)\}$, and z is the new point at which extending p is infeasible. One can verify these counterexamples using the code referenced in Section 4.1, which was used to gain intuition about them.

□

Combining Propositions 2.1, 4.6, 4.7 and 4.8 proves Theorem 4.3. This theorem shows that the function classes admitting a simple, in the sense of (4.36), tight discrete description are in fact limited to variations of smooth (weakly/strongly)-convex functions. Characterizing other interesting function classes via the constraint-based approach would thus require more general classes of constraints, e.g., constraints defined via multiple inequalities, or terwise constraints.

4.7 Concluding observations

This chapter presented a constraint-based approach to interpolation, treating constraints, rather than function or operator classes, as primary objects

of interest. Their algebraic properties determine whether they are interpolable and define meaningful classes. The chapter also introduced the notion of *pointwise extensibility* as a tool to assess interpolability, either by constructing counterexamples or by proving that a candidate constraint is indeed interpolable.

The next chapter takes a step further using pointwise extensibility and propose a principled procedure to strengthen non-interpolable constraints. This method yields improved descriptions of function or operator classes from initial (possibly loose) ones, without requiring tedious proofs of equivalence when imposed everywhere, such as the one needed, for instance, for weakly convex functions with bounded subgradients.

5

A constructive approach to interpolation conditions

This chapter results from a collaboration with Adrien Taylor.

This chapter builds on the constraint-based approach introduced in Chapter 4, and in particular on Lemma 4.2, and proposes a systematic approach to computing interpolation conditions for $\mathcal{F}_p(\mathbb{R}^d)$, starting from p itself. Specifically, given an initial description p of a given function or operator class, we constructively *strengthen* p via a refinement procedure. That is, as illustrated on Figure 5.1, we iteratively propose alternative conditions $p^{(1)}, p^{(2)}, \dots$ that are by construction equivalent to p on \mathbb{R}^d , bypassing any fastidious proof of equivalence on \mathbb{R}^d , but stronger than p on subsets $\mathcal{X} \subseteq \mathbb{R}^d$. Those alternatives become stronger and stronger as we proceed (i.e., the sets of mappings satisfying them get smaller), and ultimately converge to some limiting constraint $p^{(\infty)}$, which serves as interpolation condition for $\mathcal{F}_p(\mathbb{R}^d)$. This limiting constraint may be p itself.

The chapter is structured as follows. Section 5.2 presents the refinement procedure, and Section 5.3 comments on its implementation. Then, Section 5.4 illustrates the procedure on the classes of L -smooth convex functions and L -Lipschitz operators. Finally, Section 5.5 exploits this procedure to obtain improved conditions for function and operator classes for which no interpolation condition exists yet, including some of the classes introduced

in Section 4.4.

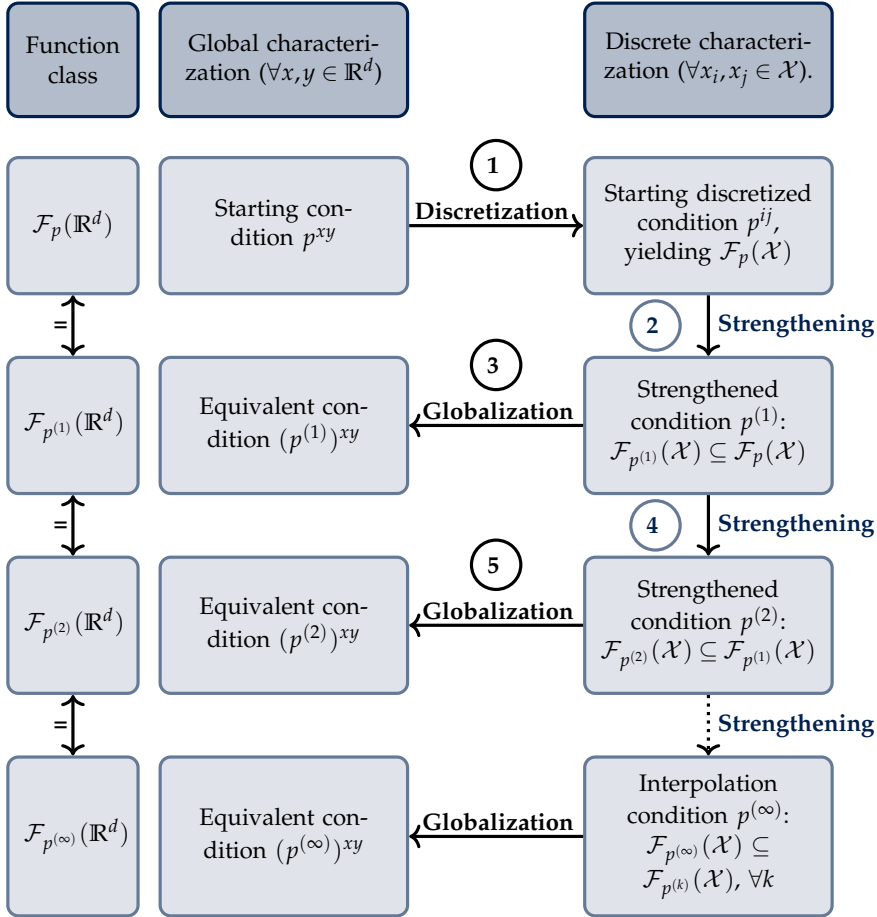


Fig. 5.1 Illustration of the iterative strengthening procedure described in Section 5.2, starting from a pairwise constraint. For an illustration of the procedure on the class $\mathcal{F}_{0,L}$ of smooth convex functions, see Section 5.4.

Notation Section 5.2 relies on cardinality of subsets \mathcal{X} of \mathbb{R}^d . To avoid treating infinite sets separately, we introduce the notation $|\mathcal{X}| = |\mathbb{N}| = \aleph_0$ for \mathcal{X} countably infinite, and $|\mathcal{X}| = |\mathbb{R}^d| = \mathfrak{c}$ for \mathcal{X} uncountable, where $K \leq \aleph_0 \leq \mathfrak{c}, \forall K \in \mathbb{N}$.

5.1 Codes

Codes for reproducing technical parts of this chapter can be found at

https://github.com/AnneRubbens/Constraints_Strengthening

5.2 Iterative strengthening of a constraint

As in Chapter 4, consider a constraint $p : (\mathbb{R}^d \times \mathbb{R}^D)^m \rightarrow \mathbb{R}^\ell$, and a domain $\bar{\mathcal{X}} \subseteq \mathbb{R}^d$. Interpolation conditions for $\mathcal{F}_p(\bar{\mathcal{X}})$ (Definition 4.3) consist of a set of constraints to impose on mappings F whose domain is a subset $\mathcal{X} \subset \bar{\mathcal{X}}$, and ensuring that such mappings admit an extension satisfying p to $\bar{\mathcal{X}}$, independently of \mathcal{X} and F . We take a first step into obtaining interpolation conditions for $\mathcal{F}_p(\bar{\mathcal{X}})$, and seek for a set of constraints (additional to p , or p itself) to impose on mappings F whose domain is a subset $\mathcal{X} \subset \bar{\mathcal{X}}$ of cardinality K , where $m \leq K \leq |\bar{\mathcal{X}}|$, and ensuring that such mappings admits an extension satisfying p to any $z \in \bar{\mathcal{X}}$.

We denote by $\tilde{\mathcal{F}}_{K,\bar{\mathcal{X}}} : (\mathbb{R}^d \times \mathbb{R}^D)^K \rightarrow \mathbb{R}^\ell$ this set of constraints, serving as a stronger definition of $\mathcal{F}_p(\bar{\mathcal{X}})$, that thus satisfies

- $\mathcal{F}_{\tilde{\mathcal{F}}_{K,\bar{\mathcal{X}}}}(\bar{\mathcal{X}}) = \mathcal{F}_p(\bar{\mathcal{X}})$,
- $\forall \mathcal{X} \subseteq \bar{\mathcal{X}}, \mathcal{F}_{\tilde{\mathcal{F}}_{K,\bar{\mathcal{X}}}}(\mathcal{X}) \subseteq \mathcal{F}_p(\mathcal{X})$, and
- Given any subset \mathcal{X} of $\bar{\mathcal{X}}$ of cardinality K , any $F \in \mathcal{F}_{\tilde{\mathcal{F}}_{K,\bar{\mathcal{X}}}}(\mathcal{X})$, and any $z \in \bar{\mathcal{X}}, \exists \bar{F} \in \mathcal{F}_p(\mathcal{X} \cup z)$ s.t. $F(x) = \bar{F}(x), \forall x \in \mathcal{X}$.

Remark 5.1. The number of points and mapping values taken as arguments by p and $\tilde{\mathcal{F}}_{K,\bar{\mathcal{X}}}$, respectively, varies from m to K .

We construct $\tilde{\mathcal{F}}_{K,\bar{\mathcal{X}}}$ as follows. By Lemma 4.2, given $\mathcal{X}' \subseteq \bar{\mathcal{X}}, |\mathcal{X}'| = K, F \in \mathcal{F}_p(\mathcal{X}')$, and $z \in \bar{\mathcal{X}}$, it is possible to extend F from \mathcal{X}' to z if and only if

$$0 \geq \min_{\substack{\tau \in \mathbb{R}, \\ F_z \in \mathbb{R}^D}} \left\{ \tau \text{ s.t. } p(\pi((x_1, F_{x_1}), \dots, (x_{m-1}, F_{x_{m-1}}), (z, F_z))) \leq \tau, \quad (5.1) \right.$$

$$\left. \begin{aligned} & \forall (x_1, x_2, \dots, x_{m-1}) \in \mathcal{X}'^{m-1}, \forall F_{x_i} \in F(x_i), i \in [m-1], \\ & \text{and for all permutation } \pi \end{aligned} \right\}.$$

Section 4.4 built on Lemma 4.2 to obtain the worst z , \mathcal{X} , and F , for which (5.1) cannot hold. This allowed assessing interpolability of constraints. We now seek to strengthen constraints which are not interpolable, by enforcing satisfaction of (5.1) for any $z \in \bar{\mathcal{X}}$, hence for the “worst” such z , i.e.,

$$0 \geq \max_{z \in \bar{\mathcal{X}}} \min_{\substack{\tau \in \mathbb{R}, \\ F_z \in \mathbb{R}^D}} \left\{ \tau \text{ s.t. } p(\pi((x_1, F_{x_1}), \dots, (x_{m-1}, F_{x_{m-1}}), (z, F_z))) \leq \tau, \quad (5.2) \right.$$

$$\left. \begin{aligned} & \forall (x_1, x_2, \dots, x_{m-1}) \in \mathcal{X}'^{m-1}, \forall F_{x_i} \in F(x_i), i \in [m-1], \\ & \text{and for all permutation } \pi \end{aligned} \right\}.$$

Let $\tilde{p}_{K, \bar{\mathcal{X}}}$ be the solution to the optimization problem defined in (5.2). Then, any $F \in \mathcal{F}_{\tilde{p}_{K, \bar{\mathcal{X}}}}(\mathcal{X}')$, hence satisfying

$$\tilde{p}_{K, \bar{\mathcal{X}}}((x_1, F_{x_1}), \dots, (x_K, F_{x_K})) \leq 0 \Leftrightarrow (5.1) \text{ is satisfied for any } z \in \bar{\mathcal{X}},$$

admits an extension in $\mathcal{F}_p(\mathcal{X}' \cup z)$, for any $z \in \bar{\mathcal{X}}$. In particular, letting $z \in \mathcal{X}'$ naturally implies $F \in \mathcal{F}_p(\mathcal{X}')$. We now formally define $\tilde{p}_{K, \bar{\mathcal{X}}}$ and prove that $\mathcal{F}_{\tilde{p}_{K, \bar{\mathcal{X}}}}(\bar{\mathcal{X}}) = \mathcal{F}_p(\bar{\mathcal{X}})$, and $\tilde{p}_{K, \bar{\mathcal{X}}}$ is stronger than p on $\bar{\mathcal{X}}$.

Definition 5.1 (*K-wise one-point strengthening*). Consider a constraint $p : (\mathbb{R}^d \times \mathbb{R}^D)^m \rightarrow \mathbb{R}$, $\bar{\mathcal{X}} \subseteq \mathbb{R}^d$, and $K \in \mathbb{N}, m \leq K \leq |\bar{\mathcal{X}}|$. We define $\tilde{p}_{K, \bar{\mathcal{X}}} : (\mathbb{R}^d \times \mathbb{R}^D)^K \rightarrow \mathbb{R}$, the *K-wise one-point strengthening* of p with respect to $\bar{\mathcal{X}}$, as

$$\begin{aligned} & \tilde{p}_{K, \bar{\mathcal{X}}}((x_1, F_{x_1}), \dots, (x_K, F_{x_K})) \quad (5.3) \\ = & \max_{z \in \bar{\mathcal{X}}} \min_{F_z \in \mathbb{R}^D, \tau \in \mathbb{R}} \left\{ \tau \text{ s.t. } p(\pi((x_1, F_{x_1}), \dots, (x_{m-1}, F_{x_{m-1}}), (z, F_z))) \leq \tau, \right. \\ & \forall \text{ tuple in } \{x_1, \dots, x_K\}^{m-1}, \forall F_{x_i} \in F(x_i), i \in [m-1], \\ & \left. \text{and for all permutation } \pi \right\}. \end{aligned}$$

Lemma 5.1 (*Properties of one-point strengthenings*). Given $\bar{\mathcal{X}} \subseteq \mathbb{R}^d$, $K \in \mathbb{N}, K \leq |\bar{\mathcal{X}}|$, consider a constraint p and its one-point strengthening $\tilde{p}_{K, \bar{\mathcal{X}}}$, as defined in (5.3). Then, $\forall \mathcal{X} \subseteq \bar{\mathcal{X}}, \mathcal{F}_{\tilde{p}_{K, \bar{\mathcal{X}}}}(\mathcal{X}) \subseteq \mathcal{F}_p(\mathcal{X})$, and $\mathcal{F}_p(\bar{\mathcal{X}}) = \mathcal{F}_{\tilde{p}_{K, \bar{\mathcal{X}}}}(\bar{\mathcal{X}})$.

Proof. Taking $z \in \mathcal{X}$ in (5.3), it holds that all $F \in \mathcal{F}_{\tilde{p}_{K, \bar{\mathcal{X}}}}(\mathcal{X})$ satisfy p on \mathcal{X} , hence the first implication holds. Similarly, all $F \in \mathcal{F}_{\tilde{p}_{K, \bar{\mathcal{X}}}}(\bar{\mathcal{X}})$ satisfy p on $\bar{\mathcal{X}}$, hence $\mathcal{F}_{\tilde{p}_{K, \bar{\mathcal{X}}}}(\bar{\mathcal{X}}) \subseteq \mathcal{F}_p(\bar{\mathcal{X}})$.

Suppose now by contradiction the existence of some $F \in \mathcal{F}_p(\bar{\mathcal{X}})$ which does not belong to $\mathcal{F}_{\tilde{p}_{K,\bar{\mathcal{X}}}}(\bar{\mathcal{X}})$. Then, at some $(x_1, \dots, x_{m-1}) \in \bar{\mathcal{X}}$, for some $F_{x_i} \in F(x_i)$, $i \in [m-1]$, and for some permutation π , it holds that

$$0 < \max_{z \in \bar{\mathcal{X}}} \min_{\tau \in \mathbb{R}, F_z \in \mathbb{R}^D} \tau \text{ s.t. } p(\pi((x_1, F_{x_1}), \dots, (x_{m-1}, F_{x_{m-1}}), (z, F_z))) \leq \tau.$$

Hence, at some $z \in \bar{\mathcal{X}}$ there exists no F_z extending p from (x_1, \dots, x_{m-1}) to z , in contradiction with the assumption $F \in \mathcal{F}_p(\bar{\mathcal{X}})$. \square

Remark 5.2. Definition 5.1 allows evaluating $\tilde{p}_{K,\bar{\mathcal{X}}}$ on any set $\mathcal{X} \subseteq \bar{\mathcal{X}}$, by simply enforcing the solution to (5.3) to be non-positive the K^{m-1} tuples in \mathcal{X}^K . Alternatively, in some cases, one can obtain a closed-form expression of $\tilde{p}_{K,\bar{\mathcal{X}}}$ by computing the solution to (5.3) in a *parametric* way in $\{x_1, \dots, x_K\}$, i.e., keeping these as parameters. Evaluating $\tilde{p}_{K,\bar{\mathcal{X}}}$ on any set $\mathcal{X} \subseteq \bar{\mathcal{X}}$ then reduces to simply imposing a set of conditions on $F: \mathcal{X} \rightarrow \mathbb{R}^D$, without having to solve an optimization problem. All examples and applications of this chapter resort to this parametric formulation of $\tilde{p}_{K,\bar{\mathcal{X}}}$.

In what follows, when $K, \bar{\mathcal{X}}$ are clear from the context, we drop the explicit dependence on $K, \bar{\mathcal{X}}$ and denote $\tilde{p}_{K,\bar{\mathcal{X}}}$ via the simpler \tilde{p} instead.

Example: convex functions

Consider the constraint $(p_{0,\infty})$, characterizing convex functions. Let $\mathcal{X} = \{(x_1, x_2)\}$, $F: \mathcal{X} \rightarrow \mathbb{R}^{1+d}: x_i \rightarrow (f_i, g_i)$, and $\bar{\mathcal{X}} = \mathbb{R}^d$. The pairwise ($K=2$) strengthening of $(p_{0,\infty})$ on \mathbb{R}^d is given by

$$\begin{aligned} \tilde{p}_{0,\infty}^{1,2} = \max_{z \in \mathbb{R}^d} \min_{\substack{\tau \in \mathbb{R}, \\ f_z \in \mathbb{R}, g_z \in \mathbb{R}^d}} \tau \text{ s.t. } & -f_z + f_i + \langle g_i, z - x_i \rangle \leq \tau, \quad i = 1, 2 \\ & -f_i + f_z + \langle g_z, x_i - z \rangle \leq \tau, \quad i = 1, 2. \end{aligned}$$

One can then verify that $\tilde{p}_{0,\infty}^{1,2} = p_{0,\infty}$.

Stability of a strengthening. Given $\bar{\mathcal{X}} \subseteq \mathbb{R}^d$, $K \in \mathbb{N}, K \leq |\bar{\mathcal{X}}|$, and starting from a given constraint p , one can strengthen p by iteratively computing one-point strengthenings $p^{(1)} = \tilde{p}_{K,\bar{\mathcal{X}}}$, $p^{(2)} = \tilde{p}_{K,\bar{\mathcal{X}}}^{(1)}$, \dots , with the goal of reaching a constraint $p_{K,\bar{\mathcal{X}}}^{(\infty)}$ that is *stable under K -wise one-point strengthening*, i.e., $p_{K,\bar{\mathcal{X}}}^{(\infty)} = \tilde{p}_{K,\bar{\mathcal{X}}}^{(\infty)}$ in the sense that $\mathcal{F}_{p_{K,\bar{\mathcal{X}}}^{(\infty)}}(\mathcal{X}) = \mathcal{F}_{\tilde{p}_{K,\bar{\mathcal{X}}}^{(\infty)}}(\mathcal{X})$, $\forall \mathcal{X} \subseteq \bar{\mathcal{X}}'$. Note

that such stable constraint might not exist.

Definition 5.2 (Stability under one-point strengthening). Given $\bar{\mathcal{X}} \subseteq \mathbb{R}^d$, $K \in \mathbb{N}, K \leq |\bar{\mathcal{X}}|$, we say a constraint p is stable under K -wise one-point strengthening with respect to $\bar{\mathcal{X}}$ if $\tilde{p}_{K, \bar{\mathcal{X}}} = p$, where $\tilde{p}_{K, \bar{\mathcal{X}}}$ is defined in (5.3).

The iterative procedure can, for instance, reach a stable constraint in finite time, see, e.g., Proposition 5.1, or as a limit to which the sequence of constraints, under their closed-form, converges, see, e.g., Proposition 5.2.

K -wise one-point extensible constraints and extension constraints. The final constraint of the refinement procedure, $p^{(\infty)}$, is K -wise pointwise extensible, in the sense that it can be extended from any subset $\mathcal{X} \subseteq \bar{\mathcal{X}}$ of cardinality K , to any $z \in \bar{\mathcal{X}}$. To obtain a *global pointwise extensible constraint* for $\bar{\mathcal{X}}$, i.e., a constraint that can be extended from any finite subset to any additional point, it then suffices to relaunch the procedure for a sequence of larger and larger K , i.e., $K_1 \leq K_2 \leq \dots \leq |\bar{\mathcal{X}}|$, with typically $K_i = K_{i-1} + 1$, to obtain a sequence $p_{K_1, \bar{\mathcal{X}}}^{(\infty)}, p_{K_2, \bar{\mathcal{X}}}^{(\infty)}, \dots, p_{K_k, \bar{\mathcal{X}}}^{(\infty)}$, where each $p_{K_k, \bar{\mathcal{X}}}^{(\infty)}$ is computed by initializing the procedure at $p_{K_i, \bar{\mathcal{X}}}^{(0)} = p_{K_{i-1}, \bar{\mathcal{X}}}^{(\infty)}$.

Under the regularity assumptions introduced in Definition 4.6, the resulting global pointwise extensible condition is an interpolation condition for $\mathcal{F}_p(\mathbb{R}^d)$, see Theorem 4.1.

Compact notation for a constraint and its strengthening. We build on the compact notation introduced in Chapter 4, and denote, for any constraint p , (i) with an index (e.g., p_α) the possible parameters for this constraint, (ii) with a superscript the points at which it is evaluated, (iii) with a tilde the corresponding one-point strengthening (5.3) of p (e.g., \tilde{p}_α the one-point strengthening of p_α obtained by (5.3)), and (iv) with a superscript in between parentheses the iterates of the one-point strengthening procedure (e.g., $p_\alpha^{(i+1)} = \tilde{p}_\alpha^{(i)}$).

Similarly to Chapter 4, we refer to constraints characterizing operators using the letter q instead of p , with all variations around this notation also used (e.g., q_α and its one-point strengthening \tilde{q}_α).

5.3 One-point strengthenings in practice

Starting from a given p , obtaining $\tilde{p}_{K,\bar{x}}$ requires solving the minimax problem (5.3). We tackle this task by dualizing the inner problem in (5.3), which can be done explicitly for a range of natural conditions used in the optimization community. Then, whenever there is no duality gap, solving (5.3) amounts to solving a single maximization problem, whose number of variables grows with ℓ , the number of inequalities in p , and K . There is no general recipe to solve this maximization problem, which often admits several solutions that can lose the initial structure of p in F .

Hence, in most cases, we do not solve (5.3) exactly but instead propose relaxations of $\tilde{p}_{K,\bar{x}}$, which might not be optimal solutions to (5.3), but are nevertheless stronger constraints than p . For similar reasons, we are, in general, not able to prove stability of a one-point strengthening, even though some insight into this question is provided on the example of Lipschitz operators (Section 5.4.2).

5.4 Examples: smooth convex functions and Lipschitz operators

Before exploiting the strengthening procedure introduced in Section 5.2 to tackle function and operator classes for which no interpolation condition is yet available, we consider two known and well-studied cases: that of L -smooth convex functions and that of L -Lipschitz operators. The first case illustrates the iterative procedure by recovering interpolation conditions from different initial descriptions of the class. The second case illustrates how showing stability of a constraint allows asserting its interpolability.

5.4.1 Recovering interpolation conditions: L -smooth convex functions

As stated in Chapter 2, the class $\mathcal{F}_{0,L}(\mathbb{R}^d)$ of L -smooth convex functions on \mathbb{R}^d can be described in numerous equivalent ways. We quickly recall three of such characterizations, equivalent when imposed everywhere. The first one, $(p_{0,L})$, consists in juxtaposing a convexity condition (2.3), with

Lipschitz continuity of the gradients (2.5).

$$p_{0,L}^{xy} \leq 0 \Leftrightarrow \begin{cases} f(x) & \geq f(y) + \langle g(y), x - y \rangle \\ \|g(x) - g(y)\|^2 & \leq L^2 \|x - y\|^2. \end{cases} \quad (p_{0,L})$$

Lipschitz continuity of the gradients is squared as it simplifies the computation of the one-point strengthening of $(p_{0,L})$. Indeed, while, in theory, one could start with both formulations, i.e., squared or not; dualizing the inner problem in (5.3) starting from the squared version, hence quadratic in the gradient, turns out way easier. Another description of $\mathcal{F}_{0,L}(\mathbb{R}^d)$, $(p'_{0,L})$, already introduced in (4.2), consists in juxtaposing convexity and another definition of smoothness, i.e., (2.6), obtained by integration of (2.5).

$$p'_{0,L} \leq 0 \Leftrightarrow \begin{cases} f(x) & \geq f(y) + \langle g(y), x - y \rangle \\ f(x) & \leq f(y) + \langle g(y), x - y \rangle + \frac{L}{2} \|x - y\|^2. \end{cases} \quad (p'_{0,L})$$

Finally, $(p''_{0,L})$, already introduced in (2.10), is an interpolation condition for $\mathcal{F}_{0,L}(\mathbb{R}^d)$, hence a third characterization of this class.

$$p''_{0,L} \leq 0 \Leftrightarrow f(x) \geq f(y) + \langle g(y), x - y \rangle + \frac{\|g(x) - g(y)\|^2}{2L}. \quad (p''_{0,L})$$

We show how to obtain the interpolation constraint for this class, $(p''_{0,L})$, first from $(p'_{0,L})$ and then from $(p_{0,L})$, without resorting to integration. To this end, consider first the following one-point strengthening.

Proposition 5.1 (Recovering $(p''_{0,L})$ from $(p'_{0,L})$). *Let $\mathcal{X} = \{(x_1, x_2)\}$ and consider $(p'_{0,L})$. Denote by $\tilde{p}'_{0,L}$ the pairwise one-point strengthening of $(p'_{0,L})$ on \mathbb{R}^d . It holds that $\tilde{p}'_{0,L} = (p''_{0,L})$.*

Proof. The one-point strengthening of $(p'_{0,L})$, $\tilde{p}'_{0,L}$, is given by

$$\tilde{p}'_{0,L} \stackrel{ij}{=} \max_{z \in \mathbb{R}^d} \min_{\substack{\tau \in \mathbb{R} \\ g_z \in \mathbb{R}^d, f_z \in \mathbb{R}}} \tau \quad (\tilde{p}'_{0,L})$$

$$\text{s.t. } f_z - f_i - \langle g_i, z - x_i \rangle - \frac{L}{2} \|z - x_i\|^2 \leq \tau, \quad i = 1, 2 \quad (\lambda_i)$$

$$f_i - f_z + \langle g_z, z - x_i \rangle - \frac{L}{2} \|z - x_i\|^2 \leq \tau, \quad i = 1, 2, \quad (\alpha_i)$$

$$f_z - f_i - \langle g_z, z - x_i \rangle \leq \tau, \quad i = 1, 2 \quad (\delta_i)$$

$$f_i - f_z + \langle g_i, z - x_i \rangle \leq \tau, \quad i = 1, 2, \quad (\mu_i),$$

where $\lambda_i, \alpha_i, \delta_i$ and μ_i are the associated dual coefficients. Setting all dual coefficients to zero except for $\lambda_i = \mu_j = \frac{1}{2}$ for some choice $i \neq j \in \{1, 2\}$, and dualizing the inner problem yields

$$\tilde{p}'_{0,L}{}^{ij} \geq \frac{1}{2} \max_{z \in \mathbb{R}^d} f_j + \langle g_j, z - x_j \rangle - f_i - \langle g_i, z - x_i \rangle - \frac{L}{2} \|z - x_i\|^2,$$

which is maximized in $z = x_i + \frac{g_i - g_j}{L}$, hence

$$0 \geq \tilde{p}'_{0,L}{}^{ij} \Leftrightarrow 0 \geq f_j - f_i + \langle g_j, x_i - x_j \rangle + \frac{1}{2L} \|g_j - g_i\|^2 \Leftrightarrow 0 \geq p''_{0,L}{}^{ij}.$$

Starting from $(p'_{0,L})$, we thus recover $(p''_{0,L})$ after a single round of one-point strengthenings. This observation was already emphasized, e.g., in [Tay17, Remark 3.15]. As a second step in our example, we show how to recover $(p''_{0,L})$ from $(p_{0,L})$, which requires several iterations of one-point strengthenings that all share the same structure. We first present a single iteration of the strengthening procedure starting with $(p_{0,L})$ (when $\alpha = 0$). □

Proposition 5.2 (One-point strengthening of p_α). *Consider $\mathcal{X} = \{(x_1, x_2)\}$, $\alpha \geq 0$, and p_α a pairwise condition defined by $\forall i, j = 1, 2$,*

$$p_\alpha^{ij} \leq 0 \Leftrightarrow \begin{cases} f_i \geq f_j + \langle g_j, x_i - x_j \rangle + \frac{\alpha}{2L} \|g_i - g_j\|^2 \\ \|g_i - g_j\|^2 \leq L^2 \|x_i - x_j\|^2. \end{cases} \quad (p_\alpha)$$

Then, \tilde{p}_α , the pairwise one-point strengthening of p_α with respect to \mathbb{R}^d , is satisfied only if, $\forall i, j = 1, 2$,

$$\tilde{p}_\alpha^{ij} \leq 0 \Rightarrow \begin{cases} f_i \geq f_j + \langle g_j, x_i - x_j \rangle + \frac{(1+\alpha)/2}{2L} \|g_i - g_j\|^2 \\ \|g_i - g_j\|^2 \leq L^2 \|x_i - x_j\|^2. \end{cases} \quad (5.4)$$

Proof. The one-point strengthening of p_α is given by

$$\begin{aligned} \tilde{p}_\alpha(f, \mathcal{X}) &= \max_{z \in \mathbb{R}^d} \min_{\substack{\tau \in \mathbb{R}, \\ g_z \in \mathbb{R}^d, f_z \in \mathbb{R}}} \tau \\ \text{s.t. } & -f_z + f_i + \langle g_i, z - x_i \rangle + \frac{\alpha}{2L} \|g_i - g_z\|^2 \leq \tau, \quad i = 1, 2 \quad (\lambda_i) \\ & -f_i + f_z - \langle g_z, z - x_i \rangle + \frac{\alpha}{2L} \|g_i - g_z\|^2 \leq \tau, \quad i = 1, 2 \quad (\mu_i) \\ & \frac{1}{2L} \|g_z - g_i\|^2 - \frac{L}{2} \|z - x_i\|^2 \leq \tau, \quad i = 1, 2, \quad (\delta_i) \end{aligned}$$

where $\lambda_i, \mu_i, \delta_i$, $i = 1, 2$, are the associated dual coefficients. Dualization of the inner problem yields

$$\begin{aligned} \tilde{p}_\alpha(f, \mathcal{X}) \geq & \max_{\substack{z \in \mathbb{R}^d, \\ \lambda_i, \mu_i, \delta_i \geq 0 \\ \sum_i (\lambda_i + \mu_i + \delta_i) = 1}} \sum_{i=1}^2 (\lambda_i (f_i + \langle g_i, z - x_i \rangle + \frac{\alpha}{2L} \|g_i - g_z^*\|^2) \\ & + \mu_i (-f_i - \langle g_z^*, z - x_i \rangle + \frac{\alpha}{2L} \|g_i - g_z^*\|^2) \\ & + \frac{\delta_i}{2L} (\|g_z^* - g_i\|^2 - L^2 \|z - x_i\|^2)) \\ \text{s.t. } 0 = & \sum_{i=1}^2 (\lambda_i - \mu_i), \\ 0 = & \sum_{i=1}^2 \left((\mu_i + \lambda_i) \frac{\alpha}{L} (g_z^* - g_i) - \mu_i (z - x_i) + \frac{\delta_i}{L} (g_z^* - g_i) \right), \end{aligned}$$

where g_z^* minimizes the inner problem in \tilde{p}_α .

In particular, setting $\lambda_i = 0, \mu_i = 0, \delta_i = \frac{1}{2}$, $i = 1, 2$, $g_z^* = \frac{g_1 + g_2}{2}$ and $z = \frac{x_1 + x_2}{2}$ yields a feasible solution of $\|g_1 - g_2\|^2 - L^2 \|x_1 - x_2\|^2$, hence \tilde{p}_α cannot be satisfied unless this quantity is negative. In addition, setting $\lambda_i = 0, \mu_j = 0, \delta_j = 0, \lambda_j = \mu_i = \delta_i = \frac{1}{3}, z = x_i + \frac{g_j - g_i}{2L}$ and $g_z^* = \frac{g_i + g_j}{2}$ yields a feasible solution of $-f_j + f_i + \langle g_i, x_j - x_i \rangle + \frac{1+\alpha}{2} \frac{1}{2L} \|g_i - g_j\|^2$, $i, j = 1, 2$. Hence, \tilde{p}_α is at least as strong as (5.4). \square

To recover $(p''_{0,L})$ from $(p_{0,L})$, it suffices to successively apply the one-point strengthening of Proposition 5.2, starting with $p^{(0)} = (p_{0,L})$, to reach $p^{(\infty)} = (p''_{0,L})$.

Proposition 5.3 (Recovering $(p''_{0,L})$ from $(p_{0,L})$). *Consider $\mathcal{X} = \{(x_1, x_2)\}$, $(p_{0,L})$, the sequence (α_k) , initiated at $\alpha_0 = 0$ and defined by*

$$\alpha_{k+1} = \frac{\alpha_k + 1}{2}, \quad (5.5)$$

and the associated sequence of conditions $(p^{(k)}) := (p_{\alpha_k})$, where p_{α_k} is defined in p_α . Then, $p^{(k+1)}$ is a relaxation of $\tilde{p}^{(k)}$ as defined in (5.4). In addition, $p^{(\infty)}$, the final result of the refinement procedure defined in Section 5.2, satisfies $p^{(\infty)} = \tilde{p}_{K, \mathcal{X}}^{(\infty)} = p_{\alpha(\infty)} = p_{\alpha=1} = (p''_{0,L})$.

Proof. By Proposition 5.2 and definition of (α_k) , the first statement holds. In addition, since (α_k) converges to $\alpha(\infty) = 1$, applying recursively Proposi-

tion 5.2 starting from $\alpha = 0$ allows iteratively strengthening p until $p^{(\infty)} = p_{\alpha^{(\infty)}} = (p''_{0,L})$, which is stable under one-point strengthening. \square

Figure 5.2 illustrates the iterative refinement procedure for $\mathcal{F}_{0,L}(\mathbb{R}^d)$.




Fig. 5.2 Given $L = 1$, $(x_1, f_1, g_1) = (0, 1, 0)$, and $x_2 = 1$, allowed regions for f_2 as a function of g_2 , according to $p^{(i)} = p_{\alpha_i}$, $i = 0, \dots, 4$, where $\alpha_0 = 0$, $\alpha_{i+1} = (1 + \alpha_i)/2$, and p_{α_i} is defined in p_{α} , and according to $p^{(\infty)} = p_{\alpha=1} = (p''_{0,L})$.

Concluding observations. In this example, the iterative strengthening procedure described in Section 5.2 converges towards the same constraint even when initialized from two different descriptions of $\mathcal{F}_{0,L}(\mathbb{R}^d)$. This holds despite relaxations in the computations of the one-point strengthenings involved. In addition, in both cases, the iterative procedure converges towards an interpolation condition for $\mathcal{F}_{0,L}(\mathbb{R}^d)$. However, the initial condition given to the procedure has a non-trivial impact on the number of iterations required to reach this condition, which varies from a single iteration to an asymptotic convergence.

5.4.2 Showing stability of a condition: L -Lipschitz operators

As a second example, consider the class $\mathcal{Q}_{-L,L}(\mathbb{R}^d)$ of Lipschitz operators Q , satisfying, $\forall x, y \in \mathbb{R}^d$, the condition (2.17), that is,

$$q_{-L,L}^{xy} \leq 0 \Leftrightarrow \|Q(x) - Q(y)\| \leq L\|x - y\|. \quad (q_{-L,L})$$

It is known [Kir34, Val43] that $q_{-L,L}$ is a global interpolation constraint for $\mathcal{Q}_{-L,L}(\mathbb{R}^d)$. This section shows how to retrieve this result via the refinement procedure described in Section 5.2. Given any $N \in \mathbb{N}$, consider $\mathcal{X} = \{(x_i)\}_{i \in [N]}$. The N -wise strengthening of $q_{-L,L}$ with respect to \mathbb{R}^d is given by

$$\tilde{q}_{-L,L} = \max_{z \in \mathbb{R}^d} \min_{\substack{\tau \in \mathbb{R}, \\ q_z \in \mathbb{R}^d}} \tau \text{ s.t. } \|q_z - t_i\|^2 - L^2\|z - x_i\|^2 \leq \tau, \quad i \in [N], \quad (\lambda_i)$$

where λ_i are the associated coefficients. Observe that this strengthening is convex in t_z and τ , and that $(q_z, \tau) = (0, \max_i \|q_i\|^2 + 1)$ satisfy Slater's condition, hence there is no duality gap for this inner problem. Dualization of the inner problem yields

$$\begin{aligned} \tilde{q}_{-L,L} &= \max_{\substack{\lambda_i \geq 0, \\ \sum_i \lambda_i = 1}} \sum_i \lambda_i (\|q_z^* - q_i\|^2 - L^2\|z^* - x_i\|^2) \\ \text{s.t. } q_z^* &= \sum_i \lambda_i q_i \text{ and } z^* = \sum_i \lambda_i x_i, \end{aligned}$$

where q_z^* minimizes the inner problem in $\tilde{q}_{-L,L}$ and z^* maximizes the resulting problem. This condition is homogeneous of degree 2 in λ_i . Hence, for any $K \geq 0$, it must hold with $\tilde{\lambda}_i = K\lambda_i$ replacing λ_i , which allows removing the condition $\sum_i \lambda_i = 1$.

The second derivative of $\tilde{q}_{-L,L}$ with respect to $\lambda_i \in [N]$ is given by

$$H_{ij} = \frac{\partial^2 \tilde{q}_{-L,L}}{\partial \lambda_i \partial \lambda_j} = \begin{cases} \|q_i - q_j\|^2 - L^2\|x_i - x_j\|^2, & i \neq j \\ 0 & \text{otherwise.} \end{cases} \quad (5.6)$$

Building on the elements of H , it holds

$$\tilde{q}_{-L,L} = \max_{\lambda_i \geq 0} \sum_{i,j \in [N], j > i} \lambda_j \lambda_i (\|q_i - q_j\|^2 - L^2\|x_i - x_j\|^2). \quad (5.7)$$

Necessarily, for $\tilde{q}_{-L,L}$ to be negative, it must hold $\|q_i - q_j\|^2 - L^2\|x_i - x_j\|^2 \leq 0, i, j \in [N]$. In addition, if this condition is satisfied, then $\tilde{q}_{-L,L} \leq 0$. Hence, $\tilde{q}_{-L,L} = q_{-L,L}$.

Concluding observations In this example, the strengthening of the initial constraint $q_{-L,L}$ was quadratic in the dual variables, allowing for a straightforward verification of its stability. In general, assessing stability is difficult given (i) the dual gap that might be non-zero, and (ii) the multiple solutions to the one-point strengthening (5.3).

5.5 Applications: improved descriptions of problem classes

This section strengthens the description of a few classical classes of functions or operators, by performing a single round of the refinement procedure. Before proceeding to the derivation of these refinements, we state some general comments on their structure and computation.

SDP-representable relaxations A particularly common scenario is that of initial constraints that are Gram representable (Definition 3.1), which, as explained in Chapter 3, allows for dimension-independent and convex PEP formulations. The resulting strengthening, in all the applications considered, can then be expressed, for some $k, n \in \mathbb{N}$, as a n -dimensional polynomials of degree k : $P(\theta) \in \mathbb{R}_k[\theta_1, \dots, \theta_n]$, whose coefficients are Gram-representable, and which must be non-positive for all θ in a given interval. While the set of non-negative polynomials is non-tractable in general, it can be relaxed in several tractable ways, e.g., by requiring polynomials to be *sum of squares* (SOS) [Hil88, Las09, Mar08], which can be expressed as an SDP. It is well-known that for the special cases ($n = 1, k \in \mathbb{N}$), ($n \in \mathbb{N}, k = 1$) and ($n = 2, k = 2$), non-negative and SOS polynomials are equivalent [Hil88]. In any case, the set of SOS polynomials provides a natural path to tractable convex relaxations of positive polynomials. Alternatively, when $n = 1$, one can use the Markov-Lukasc theorem, see, e.g., [Sze39, Theorem 1.21], to express non-negativity of polynomials as SDPs. This enables formulating one-point strengthenings of Gram-representable constraints in an SDP-representable form.

A single one-point strengthening We only perform a single iteration of the refinement procedure described in Section 5.2, as the resulting one-point strengthening involves either (i) remaining dual coefficients with no straightforward closed-form, or (ii) positive semi-definite matrices. In both cases, computing a second one-point strengthening is not straightforward. One could, in principle, relax the initial one-point strengthening as to render the computation of this second one-point strengthening easier, e.g., by fixing arbitrary values for the remaining dual coefficients.

From pairwise to terwise constraints Several initial constraints p are interpolation conditions for datasets of cardinality 2, i.e., any such dataset satisfying p is $\mathcal{F}_p(\mathbb{R}^d)$ -interpolable. But, they might be satisfied by datasets of larger cardinality, which are not consistent with any function of the class. In those cases, we compute their terwise one-point strengthenings, which thus involve three points. Whether or not the obtained terwise conditions are tight or not remains an open question. In particular, the global interpolation conditions for $\mathcal{F}_p(\mathbb{R}^d)$ might involve more than three points. Whether or not one could guess, from scratch, which interpolation conditions are pairwise and which require more points, is also an open question.

Implementation All proofs in this section are detailed in the Mathematica notebooks [Inc] provided in the code referenced in Section 5.1. In addition, the SDP formulations (5.9), (5.12) and (5.16) of the improved descriptions of (i) smooth functions satisfying a Łojasiewicz condition, (ii) block-wise smooth convex functions, and (iii) (strongly) monotone Lipschitz operators, have been added to the PEPit Python library [GMG⁺22], which allows solving PEPs on these classes straightforwardly.

5.5.1 Smooth functions satisfying a Łojasiewicz condition

This section studies the class $\mathcal{F}_{\mathbb{L},\mu,L}(\mathbb{R}^d)$ of L -smooth functions f satisfying a quadratic Łojasiewicz condition, that is satisfying, $\forall x, y \in \mathbb{R}^d$ [Loj63],

$$p_{\mathbb{L},\mu,L}^{xy} \leq 0 \Leftrightarrow \begin{cases} f(x) \leq f_* + \frac{\|\nabla f(x)\|^2}{2\mu} \\ f(x) \geq f_* \\ f(x) \geq f(y) + \frac{1}{2}\langle \nabla f(x) + \nabla f(y), x - y \rangle \\ \quad + \frac{1}{4L}\|\nabla f(x) - \nabla f(y)\|^2 - \frac{L}{4}\|x - y\|^2, \end{cases} \quad (p_{\mathbb{L},\mu,L})$$

where f_* is the minimum of f , and $0 \leq \mu \leq L$. The last condition is the strongest description of smoothness of the function (Proposition 2.1).

Motivation The interest of $\mathcal{F}_{\mathbb{L},\mu,L}(\mathbb{R}^d)$ in optimization theory, is that (i) it encompasses non-convex functions arising in machine learning, see, e.g., [KNS16, Sections 2.3, 4.1], while (ii) still providing global guarantees for stationary points [Loj63, KNS16], and yielding linear convergence rates on various first-order methods, e.g., gradient descent (GD) [Pol63], coordinate descent [KNS16, Theorem 3], stochastic gradient descent (SGD) [KNS16, Theorem 4], and, under a slightly modified condition, proximal gradient methods [BNPS17, KNS16, KJH24]. In addition, for (GD) with constant stepsizes, satisfaction of the Łojasiewicz condition is not only sufficient, but also necessary for linear convergence [AdKZ23a, Theorem 5]. However, despite $\mathcal{F}_{\mathbb{L},\mu,L}(\mathbb{R}^d)$ being a subject of extensive research, no interpolation conditions are known for this class. Hence, all convergence rates on this class are a priori non-tight.

Derivation of the one-point strengthening of $(p_{\mathbb{L},\mu,L})$ Exploiting the iterative refinement technique allows strengthening $(p_{\mathbb{L},\mu,L})$ as follows.

Proposition 5.4 (Refined description of $\mathcal{F}_{\mathbb{L},\mu,L}(\mathbb{R}^d)$). *Let $\mathcal{X} = \{(x_1, x_2, x_*)\}$. Then, $\tilde{p}_{\mathbb{L},\mu,L}^{ij}$, the pairwise one-point strengthening of $(p_{\mathbb{L},\mu,L})$ with respect to \mathbb{R}^d , is satisfied only if*

$$\tilde{p}_{\mathbb{L},\mu,L}^{ij} \leq 0 \Rightarrow \begin{cases} f_i \leq f_* + \frac{\|g_i\|^2}{2\mu} \\ f_i \geq f_* + \frac{\|g_i\|^2}{2L} \\ f_i \geq f_j + \frac{1}{2} \langle g_i + g_j, x_i - x_j \rangle + \frac{1}{4L} \|g_i - g_j\|^2 - \frac{L}{4} \|x_i - x_j\|^2 \\ \quad + \max_{0 \leq \alpha \leq \frac{2\mu}{2L+\mu}} \frac{\alpha}{(1-\alpha)(2\mu-(L+\mu)\alpha)} \left(- (L-\mu) \left(f_j - f_* + \frac{\|g_j\|^2}{2L} \right) \right. \\ \quad \left. + (1-\alpha)^2 (L+\mu) \left(f_i - f_* - \frac{\|g_i\|^2}{2L} \right) \right). \end{cases} \quad (5.8)$$

The proof of Proposition 5.4 consists in (i) formulating the pairwise one-point strengthening (5.3) of $\tilde{p}_{\mathbb{L},\mu,L}$, (ii) dualizing the inner problem in this one-point strengthening as to relax it into a maximization problem over dual variables and $z \in \mathbb{R}^d$, (iii) explicitly maximizing this relaxation in z , and (iv) computing several feasible solutions to this relaxation, by setting

some dual variables to 0 and either explicitly maximize over the remaining variables (e.g., to obtain $f_i \geq f_* + \frac{\|g_i\|^2}{2L}$), or obtain a constraint involving one of the remaining variables (e.g., α in the last condition in (5.8)).

Proof. Without loss of generality, we let $f_*, x_* = 0$. The pairwise one-point strengthening of $(p_{L,\mu,L})$ on \mathbb{R}^d is given by

$$\begin{aligned} \tilde{p}_{L,\mu,L} &= \max_{z \in \mathbb{R}^d} \min_{\substack{\tau \in \mathbb{R}, \\ g_z \in \mathbb{R}^d, f_z \in \mathbb{R}}} \tau \\ \text{s.t. } f_z - f_i + \frac{1}{2} \langle g_z + g_i, x_i - z \rangle - \frac{L\|z - x_i\|^2}{4} + \frac{\|g_z - g_i\|^2}{4L} &\leq \tau, \quad i \in \mathcal{I} \quad (\mu_i) \\ f_i - f_z + \frac{1}{2} \langle g_i + g_z, z - x_i \rangle - \frac{L\|z - x_i\|^2}{4} + \frac{\|g_z - g_i\|^2}{4L} &\leq \tau, \quad i \in \mathcal{I} \quad (\lambda_i) \\ f_z - \frac{1}{2\mu} \|g_z\|^2 &\leq \tau \quad (\alpha_1) \\ -f_z &\leq \tau, \quad (\alpha_2) \end{aligned}$$

where $\mathcal{I} = \{1, 2, *\}$, and α_1, α_2 , and $\lambda_i, \mu_i, i \in \mathcal{I}$ are the dual coefficients. Dualization of the inner problem yields

$$\begin{aligned} \tilde{p}_{L,\mu,L} &\geq \max_{\substack{z \in \mathbb{R}^d, \\ \lambda_i \geq 0, \mu_i \geq 0, \\ \alpha_1, \alpha_2 \geq 0}} \sum_{i \in \mathcal{I}} \left(\lambda_i \left(f_i + \frac{1}{2} \langle g_i + g_z^*, z - x_i \rangle - \frac{L\|z - x_i\|^2}{4} + \frac{\|g_z^* - g_i\|^2}{4L} \right) \right. \\ &\quad \left. + \mu_i \left(-f_i + \frac{1}{2} \langle g_z^* + g_i, x_i - z \rangle - \frac{L\|z - x_i\|^2}{4} + \frac{\|g_z^* - g_i\|^2}{4L} \right) \right) \\ &\quad - \alpha_1 \frac{\|g_z^*\|^2}{2\mu} \\ \text{s.t. } \frac{1}{2} &= \sum_{i \in \mathcal{I}} \lambda_i + \alpha_2 = \sum_{i \in \mathcal{I}} \mu_i + \alpha_1, \\ 0 &= \sum_{i \in \mathcal{I}} \left((\mu_i + \lambda_i) \frac{1}{2L} (g_z^* - g_i) + \frac{\lambda_i - \mu_i}{2} (z - x_i) \right) - \alpha_1 \frac{g_z^*}{\mu}, \\ 0 &\leq \sum_{i \in \mathcal{I}} \frac{\mu_i + \lambda_i}{2L} - \frac{\alpha_1}{\mu}, \end{aligned}$$

where g_z^* minimizes the inner problem in $\tilde{p}_{L,\mu,L}$, and the last condition ensures that g_z^* is a minimizer and no maximizer.

In particular, setting all dual coefficients to 0 except for $\alpha_2 = \mu_i = \frac{1}{2}$, for

a given $i \in \mathcal{I}$, $z = x_i - \frac{g_i}{L}$ and $g_z^* = 0$, yields a feasible solution of $-f_i + \frac{\|g_i\|^2}{2L}$, hence $\tilde{p}_{\mathcal{L}_{\mu,L}}$ cannot be satisfied unless this quantity is negative.

In addition, setting all dual coefficients to 0 except for $\alpha_1 \leq \frac{\mu}{2L+\mu}$, $\lambda_j = \frac{1}{2}$, $\mu_i = \frac{1}{2} - \alpha_1$, for any $i, j \in \mathcal{I}$, and maximizing over z , yields the last condition in (5.8), with $\alpha = 2\alpha_1$. Hence, $\tilde{p}_{\mathcal{L}_{\mu,L}}$ cannot be satisfied unless this condition holds. \square

Remark 5.3. Setting $\alpha = 0$ in the last condition in (5.8) allows recovering $(p_{\mathcal{L}_{\mu,L}})$, hence these conditions, although relaxations of $\tilde{p}_{\mathcal{L}_{\mu,L}}$, are at least as strong as $(p_{\mathcal{L}_{\mu,L}})$. The example illustrated in Figure 5.3 allows concluding (5.8) is strictly stronger than $(p_{\mathcal{L}_{\mu,L}})$.

Seeking Gram-representable formulation, we propose an SDP reformulation of (5.8).

Derivation of an SDP reformulation of (5.8)

Proposition 5.5 (SDP-representable refined description of $\mathcal{F}_{\mathcal{L}_{\mu,L}}(\mathbb{R}^d)$). *Consider a pair $S = \{(x_i, f_i, g_i)\}_{i=1,2}$. It holds that S satisfies (5.8) if and only (i) $f_i \leq f_* + \frac{\|g_i\|^2}{2\mu}$, $i = 1, 2$, (ii) $f_i \geq f_* + \frac{\|g_i\|^2}{2L}$, $i = 1, 2$, and (iii) $\forall i, j \in \mathcal{I}$, given*

$$\begin{aligned} A &= \left(-f_i + f_j + \frac{1}{2} \langle g_i + g_j, x_i - x_j \rangle + \frac{1}{4L} \|g_i - g_j\|^2 - \frac{L}{4} \|x_i - x_j\|^2 \right), \\ B &= (L + \mu) \left(f_i - f_* - \frac{\|g_i\|^2}{2L} \right), \\ C &= (L - \mu) \left(f_j - f_* + \frac{\|g_j\|^2}{2L} \right), \end{aligned}$$

there exists some semi-definite positive matrices

$$M^{ij} = \begin{bmatrix} M_{11} & M_{12} \\ M_{12} & M_{22} \end{bmatrix} \succcurlyeq 0, \quad \bar{M}^{ij} = \begin{bmatrix} -A(2L + \mu) & \bar{M}_{12} \\ \bar{M}_{12} & \bar{M}_{22} \end{bmatrix} \succcurlyeq 0, \quad (5.9)$$

such that

$$\begin{aligned} M_{11} - \bar{M}_{11} + \frac{4\mu}{2L + \mu} \bar{M}_{12} &= -(B - C - (L + 3\mu)A), \\ 2M_{12} - 2\bar{M}_{12} + \frac{2\mu}{2L + \mu} \bar{M}_{22} &= -(L + \mu)A + 2B, \\ M_{22} - \bar{M}_{22} &= -B. \end{aligned}$$

The proof of Proposition 5.5 is deferred to Appendix 5.A.

While (5.9) is an exact reformulation of (5.8), it remains a relaxation of $\tilde{p}_{\mathbb{L},L}$ since (5.8) is itself a relaxation of this one-point strengthening.

Comparison of $(p_{\mathbb{L},L})$ and (5.8) Figure 5.3 illustrates, on a numerical example, the difference in datasets satisfying $(p_{\mathbb{L},L})$ as compared to (5.8).




Fig. 5.3 Allowed regions for f_2 as a function of g_2 , according to (5.8) and $(p_{\mathbb{L},L})$, given $\mu = \frac{1}{2}$, $L = 1$, $(x_*, f_*, g_*) = (0, 0, 0)$, $(x_1, f_1, g_1) = (1, \frac{1}{4}, \frac{1}{2})$ and $x_2 = \frac{3}{8}$.

Figure 5.4 compares the PEP-based bounds on the performance of (GD) with constant stepsizes α , obtained by relying on $(p_{\mathbb{L},L})$ or (5.8). The difference in performance guarantees only arises from these different descriptions of $\mathcal{F}_{p_{\mathbb{L},L}}(\mathbb{R}^d)$, since both bounds are the optimal ones *given the description*. The bound obtained relying on $(p_{\mathbb{L},L})$ exactly matches the state-of-the-art one, from [AdKZ23a, Theorem 3], i.e.,

$$\begin{aligned}
 & f(x_N) - f_* \tag{5.10} \\
 & \leq \begin{cases} \left(\frac{\mu(1-L\alpha)\sqrt{-2L(L+\mu)(2-L\alpha)\mu\alpha+4L^2}}{2L+\mu} \right)^{2N} (f(x_0) - f_*) & L \in (0, \frac{1}{L}) \\ \left(\frac{2+\mu\alpha \left(1 - \frac{(2-L\alpha)(2+L\alpha)}{2+\mu\alpha} \right)}{2+\mu\alpha} \right)^N (f(x_0) - f_*) & \alpha \in [\frac{1}{L}, \frac{\sqrt{3}}{L}] \\ \left(\frac{(L\alpha-1)^2}{(L\alpha-1)^2+\mu\alpha(2-L\alpha)} \right)^N (f(x_0) - f_*) & \alpha \in (\frac{\sqrt{3}}{L}, \frac{2}{L}]. \end{cases}
 \end{aligned}$$




Fig. 5.4 Comparison of PEP-based bounds on the performance of (GD) with constant stepsizes, obtained by relying on $(p_{L_{\mu,L}})$ or the stronger (5.8), for $\mu = 0.1$, and $L = 1$. The performance measure is $f(x_N) - f_*$, and the initial condition is $f(x_0) - f_* \leq 1$. Both bounds are tight given their respective descriptions of $\mathcal{F}_{p_{L_{\mu,L}}}(\mathbb{R}^d)$. The top figure displays the convergence rate with respect to the number of iterations, for $\alpha = \frac{1}{L}$, and the bottom figure displays the convergence rate after a single iteration, for different stepsizes.

Both guarantees are linear, but the convergence is faster when considering (5.8). In addition, optimizing the stepsize with respect to both bounds

yields slightly different optimal stepsizes. In particular, the actual optimized stepsize is slightly longer than the one obtained based on the non exact characterization $(p_{\mathbb{L},L})$.

5.5.2 Convex blockwise smooth functions

This section studies the class $\mathcal{F}_{0,\mathbb{L}}(\mathbb{R}^d)$ of convex blockwise \mathbb{L} -smooth functions f , that is, functions satisfying, $\forall x, y \in \mathbb{R}^d$ [Nes12],

$$p_{0,\mathbb{L}}^{xy} \leq 0 \Leftrightarrow f(x) \geq f(y) + \langle \nabla f(y), x - y \rangle + \frac{\|\nabla_i f(x) - \nabla_i f(y)\|^2}{2L_i}, \quad \forall i \in [n], \quad (p_{0,\mathbb{L}})$$

where $n \in [d]$ and $\nabla_i f$ denotes the i^{th} block-component of the gradient of f .

Motivation The interest in $\mathcal{F}_{0,\mathbb{L}}(\mathbb{R}^d)$ arises from the fact that (i) this class encompasses many large-scale structured functions, where the gradients behave differently across blocks, and (ii) coordinate methods, extensively used and studied due to their efficiency in solving large-scale optimization problems [BT13, Nes12, Wri15], converge on this class. Block-coordinate descent, which uses, to compute its iterates, a block of the gradient as an estimator of the full gradient, has been studied under various assumptions on how to select these blocks at each iteration, e.g., in a randomized way, or in a cyclic way, see, e.g., [KHG23, AdKZ23b] for recent analysis under the PEP framework.

However, no interpolation is yet known for $\mathcal{F}_{0,\mathbb{L}}(\mathbb{R}^d)$, hence all analyses are a priori not tight. In addition, it is known that $(p_{0,\mathbb{L}})$ is an interpolation condition on \mathbb{R}^d for datasets of cardinality 2, i.e., any two-points dataset satisfying $(p_{0,\mathbb{L}})$ is consistent with a function in $\mathcal{F}_{0,\mathbb{L}}(\mathbb{R}^d)$, but not for sets of higher cardinality [KHG23].

Example 5.1 (Counterexample to sufficiency of $(p_{0,\mathbb{L}})$, from [KGHN25], Table 3). Let $n = d = 2$, $L_1 = L_2 = 1$, and

$$S = \left\{ \left(\begin{pmatrix} -1 \\ 0 \end{pmatrix}, \frac{1}{2}, \begin{pmatrix} 1 \\ 0 \end{pmatrix} \right), \left(\begin{pmatrix} 1 \\ 0 \end{pmatrix}, \frac{1}{2}, \begin{pmatrix} 1 \\ 0 \end{pmatrix} \right), \left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, 0, \begin{pmatrix} 0 \\ -1 \end{pmatrix} \right) \right\}.$$

Then, S satisfies $(p_{0,\mathbb{L}})$ while not consistent with a function in $\mathcal{F}_{0,\mathbb{L}}(\mathbb{R}^d)$, see [KGHN25], Appendix A.3, for details.

Derivation of the one-point strengthening of $(p_{0,\mathbb{L}})$ Exploiting the one-point refinement technique, we strengthen $(p_{0,\mathbb{L}})$ on sets of cardinality 3 as follows.

Proposition 5.6 (Refined description of $\mathcal{F}_{0,\mathbb{L}}(\mathbb{R}^d)$). *Let $\mathcal{X} = \{(x_1, x_2, x_3)\}$. Then, $\tilde{p}_{0,\mathbb{L}}$, the terwise one-point strengthening of $(p_{0,\mathbb{L}})$ with respect to \mathbb{R}^d , is satisfied only if*

$$\begin{aligned} \tilde{p}_{0,\mathbb{L}}^{ijk} \leq 0 \Rightarrow 0 \geq & (1 - \lambda) \left(-f_i + f_j + \langle g_j, x_i - x_j \rangle + \frac{1}{2L_m} \|g_i^{(m)} - g_j^{(m)}\|^2 \right) \\ & + \lambda \left(-f_i + f_k + \langle g_k, x_i - x_k \rangle + \frac{1}{2L_m} \|g_i^{(m)} - g_k^{(m)}\|^2 \right) \\ & + \lambda(1 - \lambda) \left(\max_l \frac{1}{2L_l} \|g_j^{(l)} - g_k^{(l)}\|^2 - \frac{1}{2L_m} \|g_j^{(m)} - g_k^{(m)}\|^2 \right), \\ & \forall m \in [n], \forall \lambda \in [0, 1], \end{aligned} \quad (5.11)$$

where $g_i^{(m)}$ denotes the m^{th} component of g_i .

Observe that the maximization over l can be easily dealt with by simply imposing (5.11) to be true for every l . The proof of Proposition 5.6 is similar to that of Proposition 5.4.

Proof. The terwise one-point strengthening of $(p_{0,\mathbb{L}})$ on \mathbb{R}^d is given by

$$\begin{aligned} \tilde{p}_{0,\mathbb{L}} = \max_{z \in \mathbb{R}^d} \min_{\substack{\tau \in \mathbb{R}, \\ g_z \in \mathbb{R}^d, f_z \in \mathbb{R}}} \tau \\ \text{s.t. } f_z - f_i + \langle g_z, x_i - z \rangle + \frac{\|g_z^{(m)} - g_i^{(m)}\|^2}{2L_m} \leq \tau, \quad i \in [3], m \in [n] \quad (\mu_{i,m}) \\ f_i - f_z + \langle g_i, z - x_i \rangle + \frac{\|g_z^{(m)} - g_i^{(m)}\|^2}{2L_m} \leq \tau, \quad i \in [3], m \in [n] \quad (\lambda_{i,m}), \end{aligned}$$

where $\lambda_{i,m}, \mu_{i,m}$, $i \in [3]$, $m \in [n]$, are the associated dual coefficients. Dualization of the inner problem yields

$$\begin{aligned} \tilde{p}_{0,\mathbb{L}} = \max_{\substack{z \in \mathbb{R}^d, \\ \mu_{i,m} \geq 0, \\ \lambda_{i,m} \geq 0}} \sum_i \sum_m \mu_{i,m} (-f_i + \langle g_z^*, x_i - z \rangle) + \frac{1}{2L_m} \|g_z^{*(m)} - g_i^{(m)}\|^2 \\ + \lambda_{i,m} (f_i + \langle g_i, z - x_i \rangle) + \frac{1}{2L_m} \|g_z^{*(m)} - g_i^{(m)}\|^2 \end{aligned}$$

$$\text{s.t. } \sum_i \sum_m \mu_{i,m} = \sum_i \sum_m \lambda_{i,m} = \frac{1}{2}, \text{ and, } \forall m \in [n],$$

$$\sum_i \left(\mu_{i,m} (x_i^{(m)} - z^{(m)}) + \mu_{i,m} \frac{g_z^{*(m)} - g_i^{(m)}}{L_m} + \lambda_{i,m} \frac{g_z^{*(m)} - g_i^{(m)}}{L_m} \right) = 0,$$

where g_z^* minimizes the inner problem in $\tilde{p}_{0,\mathbb{L}}$. In particular, setting all dual coefficients to 0 except for $\mu_{i,m} = \frac{1}{2}$, $\lambda_{k,l}$ and $\lambda_{j,l} = \frac{1}{2} - \lambda_{k,l}$ for some $i, j \in [3]$ and $m, l \in [n]$ fixed yields (5.11), with $\lambda = 2\lambda_{k,l}$, hence $\tilde{p}_{0,\mathbb{L}}$ is at least as strong as (5.11). \square

Remark 5.4. The counterexample 5.1 is discarded by (5.11) with $\lambda = \frac{1}{2}$.

Remark 5.5. Setting $\lambda = 0$ or $\lambda = 1$ in (5.11) allows recovering $(p_{0,\mathbb{L}})$, hence again this relaxation of $\tilde{p}_{0,\mathbb{L}}$ is at least as strong as $(p_{0,\mathbb{L}})$. Remark 5.4 allows concluding that (5.11) is actually strictly stronger than $(p_{0,\mathbb{L}})$.

Similarly to Section 5.5.1, seeking a formulation that involves only linear expressions in f_i 's as well as in the Gram matrix of g_i 's, x_i 's, we propose a reformulation of (5.11) as an SDP.

Derivation of an SDP reformulation of (5.11)

Proposition 5.7 (SDP-representable refined description of $\mathcal{F}_{0,\mathbb{L}}(\mathbb{R}^d)$). *Consider a triplet $S = \{(x_i, f_i, g_i)\}_{i \in [3]}$. It holds that S satisfies (5.11) if and only if $\forall m \in [n]$ and $\forall i, j, k \in [3]$, given*

$$A_{ijm} = -f_i + f_j + \langle g_j, x_i - x_j \rangle + \frac{1}{2L_m} \|g_i^{(m)} - g_j^{(m)}\|^2,$$

$$B_{jkm} = \max_l \frac{1}{2L_l} \|g_j^{(l)} - g_k^{(l)}\|^2 - \frac{1}{2L_m} \|g_j^{(m)} - g_k^{(m)}\|^2,$$

there exists some

$$M_m^{ijk} = \begin{bmatrix} -A_{ijm} & M_{12} \\ M_{12} & M_{22} \end{bmatrix} \succcurlyeq 0, \quad K_m^{ijk} \geq 0, \quad (5.12)$$

such that

$$2M_{m,12}^{ijk} + K_m^{ijk} = A_{ijm} - A_{ikm} - B_{jkm},$$

$$M_{m,22}^{ijk} - K_m^{ijk} = B_{jkm}.$$

The proof of Proposition 5.7 is deferred to Appendix 5.B.

While (5.12) is an exact reformulation of (5.11), it remains a relaxation of $\tilde{p}_{0,\mathbb{L}}$ since (5.11) is itself a relaxation of this one-point strengthening.

Comparison of $(p_{0,\mathbb{L}})$ and (5.11) Figure 5.5 illustrates on a numerical example the difference in datasets satisfying $(p_{0,\mathbb{L}})$ as compared to (5.11), the relaxation of $\tilde{p}_{0,\mathbb{L}}$.




Fig. 5.5 Allowed regions for f_3 as a function of h_3 , according to (5.11) and $(p_{0,\mathbb{L}})$, given $L_x = L_y = 1$, $(x_1, f_1, g_1) = \left(\begin{pmatrix} -1 \\ 0 \end{pmatrix}, \frac{1}{2}, \begin{pmatrix} -1 \\ 0 \end{pmatrix}\right)$, $(x_2, f_2, g_2) = \left(\begin{pmatrix} 1 \\ 0 \end{pmatrix}, \frac{1}{2}, \begin{pmatrix} 1 \\ 0 \end{pmatrix}\right)$, $x_3 = \begin{pmatrix} 0 \\ 1 \end{pmatrix}$ and $g_3 = \begin{pmatrix} 0 \\ g_3^{(2)} \end{pmatrix}$.

Figure 5.4 compares the PEP-based bounds on the performance of cyclic RCD, with constant stepsizes, obtained by relying on $(p_{0,\mathbb{L}})$ (as in [KHG23]) or (5.11). Again, the difference in performance guarantees only arises from these different descriptions of $\mathcal{F}_{p_{0,\mathbb{L}}}(\mathbb{R}^d)$. In this case, using a stronger description only yields a very small improvement as compared to the classical description. Both bounds outperform the classical hand-based bounds in this setting, e.g. from [BT13, Wri15, Nes12], as highlighted in [KHG23].

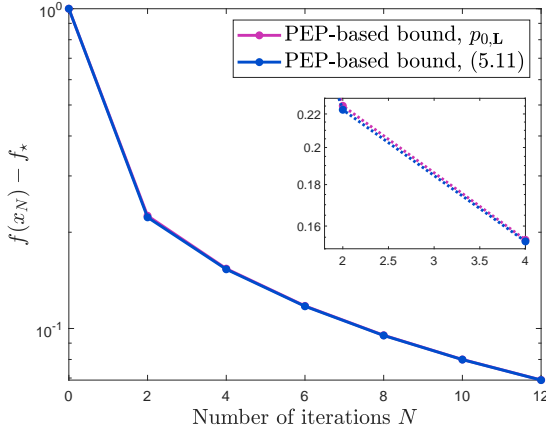


Fig. 5.6 Comparison of PEP-based bounds on the performance of cyclic RCD, with constant stepsizes $\frac{1}{L}$, obtained by relying on $(p_{0,L})$ (as in [KHG23]) or (5.11). The number of block is $n = 2$, the Lipschitz constant are $L_1 = L_2 = 1$, the performance measure is $f(x_N) - f_*$, and the initial condition is $\|x_0 - x_*\| \leq 1$.

5.5.3 Strongly monotone Lipschitz operators

This section studies the class $\mathcal{Q}_{\mu,L}(\mathbb{R}^d)$ of strongly monotone Lipschitz operators, that is operators Q satisfying, for some $0 \leq \mu \leq L$, and $\forall x, y \in \mathbb{R}^d$ [BC17, RTBG20],

$$q_{\mu,L}^{xy} \leq 0 \Leftrightarrow \begin{cases} -\langle Q(x) - Q(y), x - y \rangle + \mu \|x - y\|^2 & \leq 0 \\ \|Q(x) - Q(y)\| - L \|x - y\| & \leq 0. \end{cases} \quad (q_{\mu,L})$$

Motivation Variational inequalities of the form (VI), i.e.,

$$\text{Find } x_* \in C : 0 \in Q(x_*),$$

have been extensively studied since their introduction [Fic64, Sta64, LS67]. Indeed, many classical optimization problems can be recast as instances of (VI), e.g., problems of the form $\min_x f(x)$ for a convex function f , as

well as saddle-point problems of the form

$$\min_x \max_y \mathcal{L}(x, y),$$

for a convex-concave function $\mathcal{L}(x, y)$. Such problems naturally appear, e.g., in robust optimization [BTN02] and adversarial learning [G⁺14].

In particular, the class $\mathcal{Q}_{\mu, L}(\mathbb{R}^d)$ captures a broad range of practical problems while still sufficient to ensure existence of solutions to (VI) [BC17]. However, simple methods such as the gradient descent (GD) may diverge on $\mathcal{Q}_{\mu, L}(\mathbb{R}^d)$, hence algorithms such as the extragradient method (EG) [Kor76], whose iteration is given by

$$\begin{aligned} y_k &= x_k - \alpha Q(x_k), \\ x_{k+1} &= x_k - \alpha Q(y_k), \end{aligned} \tag{EG}$$

and the optimistic gradient descent (OG) [Pop80], with iteration

$$\begin{aligned} y_k &= x_k - \alpha Q(y_{k-1}), \\ x_{k+1} &= x_k - \alpha Q(y_k), \end{aligned} \tag{OG}$$

have been developed to address this issue. Recent work has established $\mathcal{O}(1/N)$ last-iterate convergence rates for these methods using PEP [GLG22, GTG22]. In particular, [GLG22, Theorem 3.3] proves that N iterations of (EG) with $\alpha = \frac{1}{\sqrt{2}L}$ on $\mathcal{Q}_{0, L}(\mathbb{R}^d)$ yield

$$\|Q(x_N)\|^2 \leq \frac{4L^2 \|x_0 - x_\star\|^2}{N + 1}, \tag{5.13}$$

and [GTG22, Theorem 1] proves that N iterations of (OG) with $\alpha = \frac{1}{3L}$ on $\mathcal{Q}_{0, L}(\mathbb{R}^d)$ yield

$$\|Q(x_N)\|^2 \leq \frac{123L^2 \|x_0 - x_\star\|^2}{N + 32}. \tag{5.14}$$

Nonetheless, these bounds are not a priori tight, due to the lack of interpolation conditions for $\mathcal{Q}_{\mu, L}(\mathbb{R}^d)$.

In particular, it has been shown [RTBG20, Proposition 4] that $(q_{\mu, L})$ is an interpolation condition on \mathbb{R}^d for datasets of cardinality 2, i.e., any two-points dataset satisfying $(q_{\mu, L})$ is consistent with a function in $\mathcal{Q}_{\mu, L}(\mathbb{R}^d)$, and a global interpolation condition when $d = 1$. However, when $K \geq 3$ or

$d \geq 2$, $(q_{\mu,L})$ is no longer an interpolation condition and is only necessary.

Example 5.2 (Counterexample to sufficiency of $(q_{\mu,L})$ [RTBG20]). Let $d = 2$, $\mu = 0$, $L = 1$, and

$$S = \left\{ \left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 0 \end{pmatrix} \right), \left(\begin{pmatrix} 1 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 0 \end{pmatrix} \right), \left(\begin{pmatrix} 1/2 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ -1/2 \end{pmatrix} \right) \right\}.$$

Then S satisfies $(q_{\mu,L})$ while not consistent with a function in $\mathcal{Q}_{\mu,L}$.

Derivation of the one-point strengthening of $(q_{\mu,L})$ The refinement procedure described in Section 5.2 allows strengthening $(q_{\mu,L})$, starting with a set of cardinality 3.

Proposition 5.8 (Refined description of $\mathcal{Q}_{\mu,L}(\mathbb{R}^d)$). Let $\mathcal{X} = \{(x_1, x_2, x_3)\}$. Then, $\tilde{q}_{\mu,L}$, the terwise one-point strengthening of $(q_{\mu,L})$ with respect to \mathbb{R}^d , is satisfied only if

$$\tilde{q}_{\mu,L}^{ijk} \leq 0 \Rightarrow \begin{cases} \forall \lambda_k, \mu_i, \mu_j \geq 0: \\ 0 \geq \lambda_k(\mu_i + \mu_j)(A_{ik}\mu_i + A_{jk}\mu_j) - \lambda_k\mu_i\mu_j A_{ij} \\ \quad + \lambda_k^2(B_{ik}\mu_i + B_{jk}\mu_j) + B_{ij}\mu_i\mu_j(\mu_i + \mu_j - 2\mu\lambda_k) \\ 0 \geq \lambda_k(\mu_i + \mu_j)(B_{ik}\mu_i + B_{jk}\mu_j) - \lambda_k\mu_i\mu_j B_{ij} \\ \quad + \lambda_k^2(A_{ik}\mu_i + A_{jk}\mu_j) + A_{ij}\mu_i\mu_j(\mu_i + \mu_j - 2\mu\lambda_k), \end{cases} \quad (5.15)$$

where $A_{ij} = \|t_i - t_j\|^2 - L^2\|x_i - x_j\|^2$ and $B_{ij} = -2L\langle t_i - t_j, x_i - x_j \rangle + 2L\mu\|x_i - x_j\|^2$.

The proof of Proposition 5.8 follows that of Propositions 5.4 and 5.6.

Proof. The terwise one-point strengthening of $(q_{\mu,L})$ on \mathbb{R}^d is given by

$$\tilde{q}_{\mu,L} = \max_{z \in \mathbb{R}^d} \min_{\substack{\tau \in \mathbb{R}, \\ t_z \in \mathbb{R}^d}} \tau \text{ s.t. } \frac{\|t_z - t_i\|^2}{2L} - \frac{L}{2}\|z - x_i\|^2 \leq \tau, \quad i \in [3] \quad (\lambda_i)$$

$$- \langle t_z - t_i, z - x_i \rangle + \mu\|z - x_i\|^2 \leq \tau, \quad i \in [3] \quad (\mu_i)$$

where λ_i, μ_i are the associated dual coefficients. Dualization of the inner problem yields

$$\tilde{q}_{\mu,L} \geq \frac{1}{2L} \max_{\substack{\lambda_i \geq 0, \\ \mu_i \geq 0}} \sum_i (-\lambda_i (\|t_z^* - t_i\|^2 - L^2\|z - x_i\|^2))$$

$$\begin{aligned}
 & + \mu_i 2L(-\langle t_z^* - t_i, z - x_i \rangle + \mu \|z - x_i\|^2) \\
 \text{s.t. } & \sum_i \frac{1}{L} \lambda_i (t_z^* - t_i) - \mu_i (z - x_i) = 0, \quad \sum_i (\lambda_i + \mu_i) = 1,
 \end{aligned}$$

where t_z^* minimizes the inner problem in $\tilde{q}_{\mu,L}$. Minimizing over z , multiplying by $2L((\sum_i \lambda_i)^2 + (\sum_i \mu_i)^2)$ and rearranging the terms yields $\tilde{q}_{\mu,L} \geq$

$$\begin{aligned}
 \max_{\substack{\lambda_i \geq 0, \mu_i \geq 0, \\ \sum_i (\lambda_i + \mu_i) = 1}} & \sum_{i,j \in [3]: j > i} \left(A_{ij} (\lambda_i \lambda_j \sum_k \lambda_k + \lambda_i \mu_j (\sum_k \mu_k) \right. \\
 & + \lambda_j \mu_i (\sum_k \mu_k) - (\sum_k \lambda_k) \mu_i \mu_j - 2\mu (\sum_k \mu_k) \lambda_i \lambda_j) \\
 & + B_{ij} (\mu_i \mu_j \sum_k \mu_k + \mu_i \lambda_j (\sum_k \lambda_k) \\
 & \left. + \mu_j \lambda_i (\sum_k \lambda_k) - (\sum_k \mu_k) \lambda_i \lambda_j - 2\mu (\sum_k \lambda_k) \mu_i \mu_j) \right).
 \end{aligned}$$

This condition is homogeneous of degree 3 in λ_k, μ_i, μ_j . Hence, for any $K \geq 0$, it must hold with $\tilde{\mu}_i = K\mu_i, \tilde{\lambda}_i = K\lambda_i$ replacing μ_i, λ_i respectively, which allows removing the condition $\sum_i (\lambda_i + \mu_i) = 1$.

Setting only λ_k, μ_i and μ_j non-zero yields the first condition in (5.15), and setting only μ_k, λ_i and λ_j non-zero yields the second condition in (5.15), hence $\tilde{q}_{\mu,L}$ is at least as strong as (5.15). Those choices of coefficients allow for the terms multiplying A_{ij} or B_{ij} to be potentially negative. \square

Remark 5.6. Observe that the counterexample 5.2 does not satisfy (5.15), and separates $\mathcal{Q}_{q_{\mu,L}}(\mathcal{X})$ from $\mathcal{Q}_{\tilde{q}_{\mu,L}}(\mathcal{X})$.

Remark 5.7. Setting $\lambda_k = 0$ in (5.15) allows recovering $(q_{\mu,L})$, hence again this relaxation of $\tilde{q}_{\mu,L}$ is at least as strong as $(q_{\mu,L})$. Remark 5.6 allows concluding (5.15) is actually strictly stronger than $(q_{\mu,L})$.

Similarly to Sections 5.5.1 and 5.5.2, we propose a relaxation of (5.15) as an SDP.

Derivation of an SDP reformulation of (5.15).

Proposition 5.9 (SDP-representable refined description of $\mathcal{Q}_{\mu,L}(\mathbb{R}^d)$). *With the same notation as in Proposition 5.8 and considering a triplet $S = \{(x_i, t_i)\}_{i \in [3]}$, it holds that S satisfies (5.15) if $\forall i, j, k \in [3]$, there exists some $M_{A,B}^{ijk} \succcurlyeq 0, M_{B,A}^{ijk} \succcurlyeq$*

0, where

$$M_{\alpha,\beta}^{ijk} = \begin{bmatrix} -\beta_{ij} & 0 & 0 & M_{14} & M_{15} & M_{16} & M_{17} \\ 0 & -\alpha_{jk} & -M_{15} & 0 & -M_{14} & M_{26} & M_{27} \\ 0 & -M_{15} & -\beta_{ij} & M_{34} & -M_{16} & 0 & M_{37} \\ M_{14} & 0 & M_{34} & -\beta_{jk} & -M_{27} & M_{46} & 0 \\ M_{15} & -M_{14} & -M_{16} & -M_{27} & M_{55} & -M_{37} & -M_{46} \\ M_{16} & M_{26} & 0 & M_{46} & -M_{37} & -\alpha_{ik} & 0 \\ M_{17} & M_{27} & M_{37} & 0 & -M_{46} & 0 & -\beta_{ik} \end{bmatrix}, \quad (5.16)$$

and

$$M_{55} + 2M_{17} + 2M_{26} + 2M_{34} = \alpha_{ij} + 2\mu\beta_{ij} - \alpha_{jk} - \alpha_{ik}.$$

The proof of Proposition 5.9 is deferred to Appendix 5.C.

Comparison of $(q_{\mu,L})$ and (5.16) Figure 5.7 illustrates on a numerical example the difference in datasets satisfying $(q_{\mu,L})$ as compared to (5.16), the relaxation of $\tilde{q}_{\mu,L}$.




Fig. 5.7 Allowed regions for t_3 , according to $(q_{\mu,L})$ or (5.15), given $\mu = 0$, $L = 1$, $(x_1, t_1) = \begin{pmatrix} 0 \\ 0 \end{pmatrix}$, $(x_2, t_2) = \begin{pmatrix} 1 \\ 0 \\ -1/2 \end{pmatrix}$, and $x_3 = \begin{pmatrix} 1/2 \\ 1 \end{pmatrix}$.

Figure 5.8 compares the bounds on the performance of (EG) and (OG), obtained (i) in [GLG22, Theorem 3.3] and [GTG22, Theorem 1], i.e., in a computer-aided way, relying on $(q_{\mu,L})$ and using relaxations, (ii) in a PEP-based way, using $(q_{\mu,L})$, and (iii) in a PEP-based way, using (5.16).




Fig. 5.8 Comparison of PEP-based bounds on the performance of (EG), with stepsize $\alpha = \frac{1}{\sqrt{2}L}$ (top figure) and (OG), with stepsize $\alpha = \frac{1}{3L}$ (bottom figure), obtained (i) in [GLG22, Theorem 3.3] and [GTG22, Theorem 1], i.e., in a non tight computer-aided way, relying on $(q_{\mu,L})$, (ii) in a PEP-based way, using $(q_{\mu,L})$, and (iii) in a PEP-based way, using (5.16). The performance measure is $\|Q(x_N)\|^2$, the initial condition is $\|x_0 - x_\star\|^2 \leq 1$, $\mu = 0$ and $L = 1$.

The tight PEP-based bounds improve on [GLG22, Theorem 3.3] and [GTG22, Theorem 1], and relying on (5.15) yields better results than using $(q_{\mu,L})$, from 2 iterations.

Remark 5.8. Since $(q_{\mu,L})$ is an interpolation condition on \mathbb{R}^d for datasets of cardinality 2 [RTBG20, Proposition 4], all methods achieving linear con-

vergence rates on $\mathcal{Q}_{\mu,L}$ should be analyzed using this constraint. Indeed, analyzing a single iteration of such methods, involving 2 points only, is sufficient for a global analysis. Hence, using $(p_{\mu,L})$ already yields tight guarantees.

5.5.4 Uniformly convex functions

This section studies the class $\mathcal{F}_{\mu,q}(\mathbb{R}^d)$ of uniformly convex functions f satisfying, $\forall x, y \in \mathbb{R}^d$, $(p_{\mu,q})$, that we recall here.

$$p_{\mu,q}^{xy} \leq 0 \Leftrightarrow f(x) \geq f(y) + \langle \nabla f(y), x - y \rangle + \frac{\mu}{q} \|x - y\|^q, \text{ where } \mu \geq 0, q \geq 2, \quad (p_{\mu,q})$$

Motivation The concept of uniform convexity, originally introduced in 1936 [Cla36], was then adopted in convex optimization to strengthen convergence results as compared to convex functions, see, e.g., [JN14a, JN14b, DN21]. In particular, uniformly convex functions admit a unique minimizer, and their dual functions are Hölder smooth functions, see, e.g., [Nes15, Lemma 1].

Uniformly convex functions are also curious from an interpolation point of view. Indeed, $(p_{\mu,q})$ is algebraically simple, and is an interpolation constraint for its limit case $q = 2$. However, for $q > 2$, it is not a global interpolation constraint anymore, see Section 4.4.

Extensibility from 2-points datasets We first prove that $(p_{\mu,q})$ can be extended from any set of cardinality 2 to any $z \in \mathbb{R}^d$, hence the strengthening technique proposed in Section 5.2 can only yield a stronger constraint starting from sets of cardinality ≥ 3 .

Proposition 5.10 (Extensibility from 2-points datasets of $(p_{\mu,q})$). *Let $S = \{(x_i, f_i, g_i)\}_{i=1,2}$ be a set satisfying $(p_{\mu,q})$. Given any $z \in \mathbb{R}^d$, there exist some $f_z \in \mathbb{R}$, $g_z \in \mathbb{R}^d$ such that $S \cup (z, f_z, g_z)$ satisfies $(p_{\mu,q})$.*

Proof. We seek for some $f_z \in \mathbb{R}$, $g_z \in \mathbb{R}^d$ extending $(p_{\mu,q})$ to x , that is such that $\forall i, j = 1, 2$, it holds

$$f_z \geq f_i + \langle g_i, x - x_i \rangle + \frac{\mu}{q} \|x - x_i\|^q, \quad (5.17)$$

$$f_j \geq f_z + \langle g_z, x_j - x \rangle + \frac{\mu}{q} \|x - x_j\|^q. \quad (5.18)$$

Let

$$\begin{aligned} i_* &= \operatorname{argmax}_{i=1,2} (f_i + \langle g_i, x - x_i \rangle + \frac{\mu}{q} \|x - x_i\|^q), \\ f_z &= f_{i_*} + \langle g_{i_*}, x - x_{i_*} \rangle + \frac{\mu}{q} \|x - x_{i_*}\|^q, \\ g_z &= g_{i_*} + \mu \|x - x_{i_*}\|^{q-2} (x - x_{i_*}) + K(x - x_{i_*}), \end{aligned}$$

where $K \geq 0$. Then, (5.17) is automatically satisfied. In addition, it holds

$$\begin{aligned} f_z + \langle g_z, x_{i_*} - x \rangle + \frac{\mu}{q} \|x - x_{i_*}\|^q &= f_{i_*} - \mu \|x - x_{i_*}\|^q + 2\frac{\mu}{q} \|x - x_{i_*}\|^q - K \|x - x_{i_*}\| \\ &\leq f_{i_*}, \end{aligned}$$

and

$$\begin{aligned} &f_z + \langle g_z x_j - x \rangle + \frac{\mu}{q} \|x - x_j\|^q \\ &= f_{i_*} + \langle g_{i_*}, x_j - x_{i_*} \rangle + \langle K(x - x_{i_*}), x_j - x \rangle \\ &+ \mu \|x - x_{i_*}\|^{q-2} \langle x - x_{i_*}, x_j - x \rangle + \frac{\mu}{q} (\|x - x_{i_*}\|^q + \|x - x_j\|^q) \\ &\leq f_j - \langle K X_{i_*}, X_j \rangle + \frac{\mu}{q} (\|X_{i_*}\|^q + \|X_j\|^q - \|X_{i_*} - X_j\|^q - q \|X_{i_*}\|^{q-2} \langle X_{i_*}, X_j \rangle), \end{aligned}$$

where $X_{i_*} = x - x_{i_*}$ and $X_j = x - x_j$. Suppose first $\langle X_{i_*}, X_j \rangle \leq 0$, and let $K = 0$. Then, by Bernoulli's inequality, and since $(x + y)^q \geq x^q + y^q$ for $q \geq 2$, $x, y \geq 0$, it holds

$$\begin{aligned} \|X_{i_*} - X_j\|^q &= (\|X_{i_*}\|^2 + \|X_j\|^2 - 2\langle X_{i_*}, X_j \rangle)^{q/2} \\ &= (\|X_{i_*}\|^2 + \|X_j\|^2)^{q/2} (1 - 2 \frac{\langle X_{i_*}, X_j \rangle}{\|X_{i_*}\|^2 + \|X_j\|^2})^{q/2} \\ &\geq \|X_{i_*}\|^q + \|X_j\|^q - q \langle X_{i_*}, X_j \rangle (\|X_{i_*}\|^2 + \|X_j\|^2)^{q/2-1} \\ &\geq \|X_{i_*}\|^q + \|X_j\|^q - q \langle X_{i_*}, X_j \rangle \|X_{i_*}\|^{q-2}, \end{aligned}$$

hence (5.18) holds. Suppose now $\langle X_{i_*}, X_j \rangle > 0$. Then there exists a sufficiently large K such that (5.18) holds. \square

One-point strengthening of $(p_{\mu,q})$ We now show that $(p_{\mu,q})$ can be strengthened by considering sets of higher cardinality.

Proposition 5.11 (Refined description of $\mathcal{F}_{\mu,q}(\mathbb{R}^d)$). *Let $\mathcal{X} = \{(x_1, x_2, x_3)\}$. Then, $\tilde{p}_{\mu,q}$, the terwise one-point strengthening of $(p_{\mu,q})$ on \mathbb{R}^d , is satisfied only if*

$$\begin{aligned} \tilde{p}_{\mu,q}^{ijk} \leq 0 \Rightarrow 0 \geq & \max_{\lambda \in [0,1]} \lambda(-f_i + f_k + \langle g_k, x_i - x_k \rangle) \\ & + (1 - \lambda)(-f_j + f_k + \langle g_k, x_j - x_k \rangle) \\ & + \frac{\mu}{q} \left((\lambda(1 - \lambda)^q + (1 - \lambda)\lambda^q) \|x_i - x_j\|^q \right. \\ & \left. + \|x_k - \lambda x_i - (1 - \lambda)x_j\|^q \right). \end{aligned} \quad (5.19)$$

Proof. The terwise one-point strengthening of $(p_{\mu,q})$ on \mathbb{R}^d is given by

$$\begin{aligned} \tilde{p}_{\mu,q} = & \max_{z \in \mathbb{R}^d} \min_{\substack{\tau \in \mathbb{R}, \\ g_z \in \mathbb{R}^d, f_z \in \mathbb{R}}} \tau \\ \text{s.t. } & -f_z + f_i + \langle g_i, z - x_i \rangle + \frac{\mu}{q} \|z - x_i\|^q \leq \tau, \quad i \in [3] \quad (\lambda_i) \\ & -f_i + f_z - \langle g_z, z - x_i \rangle + \frac{\mu}{q} \|z - x_i\|^q \leq \tau, \quad i \in [3] \quad (\mu_i) \end{aligned}$$

where λ_i, μ_i ($i \in [3]$) are the associated dual coefficients. Dualization of the inner problem yields

$$\begin{aligned} \tilde{p}_{\mu,q} = & \max_{\substack{z \in \mathbb{R}^d, \\ \lambda_i, \mu_i \geq 0}} \sum_{i=1}^3 (\lambda_i(f_i + \langle g_i, z - x_i \rangle + \frac{\mu}{q} \|z - x_i\|^q) \\ & + \mu_i(-f_i - \langle g_z, z - x_i \rangle + \frac{\mu}{q} \|z - x_i\|^q)) \\ \text{s.t. } & \sum_{i=1}^3 \lambda_i = \sum_{i=1}^3 \mu_i = \frac{1}{2}, \quad z = 2 \sum_{i=1}^3 (\mu_i x_i). \end{aligned}$$

In particular, setting only $\lambda_k = \frac{1}{2}$, $\mu_i \in [0, \frac{1}{2}]$, $\mu_j = \frac{1}{2} - \mu_i$ non-zero yields a feasible solution of

$$\begin{aligned} 0 \geq & \mu_i(-f_i + f_k + \langle g_k, x_i - x_k \rangle) + (\frac{1}{2} - \mu_i)(-f_j + f_k + \langle g_k, x_j - x_k \rangle) \\ & + \frac{\mu}{q} \left((\mu_i(\frac{1}{2} - \mu_i)^q + (\frac{1}{2} - \mu_i)\mu_i^q) \|x_j - x_i\|^q + \|x_k - \mu_i x_i - (\frac{1}{2} - \mu_i)x_j\|^q \right). \end{aligned}$$

Hence $\tilde{p}_{\mu,L}$ is at least as strong as (5.19). □

Remark 5.9. Setting $\lambda = 1$ in (5.19) allows recovering $(p_{\mu,q})$.

Remark 5.10. The case $p = 2$ can be simplified as

$$\begin{aligned} \tilde{p}_{\mu,2}^{ijk} \leq 0 &\Rightarrow 0 \geq \lambda(-f_i + f_k + \langle g_k, x_i - x_k \rangle + \frac{\mu}{2} \|x_i - x_k\|^2) \\ &\quad + (1 - \lambda)(-f_j + f_k + \langle g_k, x_j - x_k \rangle + \frac{\mu}{2} \|x_i - x_k\|^2) \\ &\Leftrightarrow f_i \geq f_k + \langle g_k, x_i - x_k \rangle + \frac{\mu}{2} \|x_i - x_k\|^2 \text{ and} \\ &\quad f_j \geq f_k + \langle g_k, x_j - x_k \rangle + \frac{\mu}{2} \|x_j - x_k\|^2. \end{aligned}$$

We recover the classical definition of μ -strongly convex functions.

5.5.5 Smooth convex functions with constrained convex domain

Given an open convex subset $\tilde{\mathcal{X}} \subseteq \mathbb{R}^d$, consider the class $\mathcal{F}_{0,L}(\tilde{\mathcal{X}})$ of functions that are smooth and convex on $\tilde{\mathcal{X}}$, in the sense of $(p_{0,L})$, but not necessarily outside of $\tilde{\mathcal{X}}$. It is known that $(p'_{0,L})$ no longer holds in general for this function class [Dro20, Section 2], but that it remains valid for pairs (x, y) of points $x, y \in \tilde{\mathcal{X}}$ that are sufficiently close. In particular, it holds at every $x, y \in \tilde{\mathcal{X}}$ such that [Dro20, Theorem 3.1]

$$\|x - y\| < \text{dist}(y, \mathbb{R}^d \setminus \tilde{\mathcal{X}}), \quad (5.20)$$

This result is obtained by showing that in this case, all quantities involved in the strengthening of $(p'_{0,L})$ into $(p''_{0,L})$ belong to $\tilde{\mathcal{X}}$. Building on the iterative strengthening of $(p_{0,L})$, we improve this description by showing $(p''_{0,L})$ holds at pairs that are twice as distant.

Proposition 5.12. *Let $\tilde{\mathcal{X}} \subseteq \mathbb{R}^d$ an open convex set and $f \in \mathcal{F}_{0,L}(\tilde{\mathcal{X}})$. Then, for any $x, y \in \tilde{\mathcal{X}}$ such that*

$$\|x - y\| < 2 \text{dist}(y, \mathbb{R}^d \setminus \tilde{\mathcal{X}}), \quad (5.21)$$

it holds

$$f(y) \geq f(x) + \langle \nabla f(x), y - x \rangle + \frac{1}{2L} \|\nabla f(x) - \nabla f(y)\|^2.$$

Proof. Let $z = y + \frac{g(x) - g(y)}{2L}$. It holds

$$\|z - y\| = \frac{1}{2L} \|g(y) - g(x)\| \leq \frac{1}{2} \|y - x\| \stackrel{(5.21)}{<} \text{dist}(y, \mathbb{R}^d \setminus \bar{\mathcal{X}}),$$

hence $z \in \bar{\mathcal{X}}$. Setting $x_2 = x$ and $x_1 = y$ in the proof of Proposition 5.2, observe that the strengthening of $(p_{0,L})$ into $(p'_{0,L})$ remains valid since it only involves quantities belonging to $\bar{\mathcal{X}}$. \square

Remark 5.11. This factor 2 cannot be improved via the one-point strengthening procedure, as long as one only uses $(p_{0,L})$ or $(p'_{0,L})$. However, it might be improved by starting from other conditions.

Building on [Dro20, Theorem 3.1], the author of [Dro20] obtained properties valid for all $x, y \in \bar{\mathcal{X}}$, however far from each other, that are expressed as optimization problems. Relying on Proposition 5.12 instead of [Dro20, Theorem 3.1] allows improving this description of $\mathcal{F}_{0,L}(\bar{\mathcal{X}})$.

Corollary 5.1 (From [Dro20], Corollary 3.2). *Given $\{(x, g_x, f_x), (y, g_y, f_y)\} \in (\mathbb{R}^d \times \mathbb{R}^d \times \mathbb{R})^2$, let*

$$\begin{aligned} \mathcal{U}_N(\text{resp. } \mathcal{B}_N)(x, g_x, f_x, y, g_y) &= \max(\text{resp. } \min)_{g_i \in \mathbb{R}^d, f_i \in \mathbb{R}} f_N \\ \text{s.t. } f_0 &= f_x, g_0 = g_x, g_N = g_y, \text{ and } \forall 0 \leq i < N, \\ \frac{1}{2L} \|g_i - g_{i+1}\|^2 &\leq f_i - f_{i+1} - \frac{1}{N} \langle g_{i+1}, x - y \rangle, \\ \frac{1}{2L} \|g_i - g_{i+1}\|^2 &\leq f_{i+1} - f_i - \frac{1}{N} \langle g_i, y - x \rangle, \end{aligned}$$

Then,

1. If $f \in \mathcal{F}_{0,L}(\bar{\mathcal{X}})$, then for any $x, y \in \bar{\mathcal{X}}$, and $\forall N > \frac{\|x - y\|}{\min(\text{dist}(x, \mathbb{R}^d \setminus \bar{\mathcal{X}}), \text{dist}(y, \mathbb{R}^d \setminus \bar{\mathcal{X}}))}$,

$$\mathcal{B}_N(x, \nabla f(x), f(x), y, \nabla f(y)) \leq f(y) \leq \mathcal{U}_N(x, \nabla f(x), f(x), y, \nabla f(y)).$$

2. If $\{(x, g_x, f_x), (y, g_y, f_y)\} \in (\mathbb{R}^d \times \mathbb{R}^d \times \mathbb{R})^2$ satisfy, for some $L, N > 0$,

$$\mathcal{B}_N(x, g_x, f_x, y, g_y) \leq f_y \leq \mathcal{U}_N(x, g_x, f_x, y, g_y),$$

then there exists a function $f \in \mathcal{F}_{0,L}(\bar{\mathcal{X}})$ such that $f(x) = f_x$, $f(y) = f_y$, $\nabla f(x) = g_x$, and $\nabla f(y) = g_y$.

Proof. The proof follows the same lines as in [Dro20, Corollary 3.2], replacing Proposition [Dro20, Theorem 3.1] by Proposition 5.12. \square

In [Dro20, Corollary 3.2], a similar result is proved, except for the factor 2 in the definition of N : Corollary 5.1 allows bounding $f(y)$ by $\mathcal{U}_N, \mathcal{B}_N$ for twice smaller values of N than [Dro20, Corollary 3.2].

To illustrate how this strengthens the description of $\mathcal{F}_{0,L}(\bar{\mathcal{X}})$, consider the following numerical example, from [Dro20, Figure 3]. Given $f: \bar{\mathcal{X}} \rightarrow \mathbb{R}$, suppose $L = 1$, $x = 0$, $\nabla f(x) = 0$, $f(x) = 0$, $\|y\|^2 = 1$, $\|\nabla f(y)\|^2 = \frac{1}{2}$ and $\min(\text{dist}(x, \mathbb{R}^d \setminus \bar{\mathcal{X}}), \text{dist}(y, \mathbb{R}^d \setminus \bar{\mathcal{X}})) = 1$. We show in Figure 5.9 the allowed regions for $f(y)$ as a function of $\langle \nabla f(y), y \rangle$, according to [Dro20, Corollary 3.2] and Corollary 5.1. Obtaining closed-form interpolation con-




Fig. 5.9 From [Dro20, Figure 3]. Allowed regions for $f(y)$ as a function of $\langle \nabla f(y), y \rangle$, according to [Dro20, Corollary 3.2] and Corollary 5.1, and given $L = 1$, $(x, f(x), \nabla f(x)) = (0, 0, 0)$, $\|y\|^2 = 1$, $\|\nabla f(y)\|^2 = \frac{1}{2}$ and $\min(\text{dist}(x, \mathbb{R}^d \setminus \bar{\mathcal{X}}), \text{dist}(y, \mathbb{R}^d \setminus \bar{\mathcal{X}})) = 1$. Corollary 5.1 allows bounding f_y by $\mathcal{B}_1, \mathcal{U}_1$, while [Dro20, Corollary 3.2] can only rely on $\mathcal{B}_2, \mathcal{U}_2$.

ditions for function classes with constrained domain would require, e.g., considering more specific domains as the unit ball.

5.6 Concluding observations

Chapters 4 and 5 presented a constraint-based approach to interpolation, that partially answered some of the open questions regarding interpolation, discussed in Chapter 2.

Question 1. Given a function/operator class, how to guess a candidate interpolation condition?

Question 2. How to assess a constraint's interpolability?

The notion of pointwise extensibility (Definition 4.5) furnishes a possible answer to Question 2. Indeed, (i) it allows obtaining counterexamples to a constraint's interpolability in an automated way (Section 4.4), and (ii) it reduces proving interpolability of a candidate constraint to an algebraic question, independent from the analytical properties of the class analyzed.

In addition, the one-point strengthening procedure (Definition 5.1) introduced in Chapter 5 can be seen as a principled and constructive approach to answering Question 1. Indeed, starting from any characterization of a class of functions/operators, it iteratively constructs refined characterizations of this class, obtaining interpolation conditions as a final product. In particular, it allows for strengthening given sets of conditions based only on their algebraic properties, leading to improved conditions that would have been challenging to guess based solely on the class' analytical properties. Proving stability (Definition 5.2) of constraint via the strengthening procedure also answers Question 2, but we have seen that this is in practice challenging.

Finally, the proposed constraint-based approach allows bypassing Question 1, by directly highlighting function classes based on the fact that they possess simple interpolation conditions rather than on their analytical properties.

On the downside, the constraint-based approach naturally showcases some limitations. For instance, our main technical tools for obtaining one-point strengthenings consist of dualizing the inner problem in (5.3). Hence, the approach fails as soon as the initial constraint under consideration is not convex in F_z . Similarly, the approach is naturally limited to constraints for which this dualization yields a zero duality gap. Furthermore, we quickly arrive at formulations that turn out to be hardly manipulable, both for computing their one-point strengthening and for proving their stability.

Appendix to Chapter 5

5.A Proof of Proposition 5.5

Proof. The first conditions are strictly equivalent to the first ones in (5.8). The last inequality in (5.8) can be equivalently written as

$$0 \geq A + \frac{\alpha}{(1-\alpha)(2\mu - (L+\mu)\alpha)} \left((1-\alpha)^2 B - C \right), \quad \forall \alpha \in \left[0, \frac{2\mu}{2L+\mu} \right]. \quad (5.22)$$

In the interval $\left[0, \frac{2\mu}{2L+\mu} \right]$, $\alpha \leq \frac{2\mu}{L+\mu} \leq 1$, hence (5.22) is equivalent to

$$\begin{aligned} 0 &\geq (1-\alpha)(2\mu - (L+\mu)\alpha)A + (1-\alpha)^2\alpha B - \alpha C, \\ \Leftrightarrow 0 &\geq 2\mu A + \alpha(B - C - (L+3\mu)A) + \alpha^2((L+\mu)A - 2B) + \alpha^3 B, \\ \Leftrightarrow 0 &\geq a_0 + a_1\alpha + a_2\alpha^2 + a_3\alpha^3 := P(\alpha), \end{aligned}$$

where $a_0 = 2\mu A$, $a_1 = B - C - (L+3\mu)A$, $a_2 = (L+\mu)A - 2B$ and $a_3 = B$. By Markov-Lukacs theorem [Sze39, Theorem 1.21],

$$\begin{aligned} P(\alpha) \leq 0, \quad \forall \alpha \in \left[0, \frac{2\mu}{2L+\mu} \right] \\ \Leftrightarrow \exists M, \bar{M} \in \mathbb{S}_+^2 \text{ s.t. } P(\alpha) = -\alpha \begin{bmatrix} 1 & \alpha \end{bmatrix} M \begin{bmatrix} 1 \\ \alpha \end{bmatrix} - \left(\frac{2\mu}{2L+\mu} - \alpha \right) \begin{bmatrix} 1 & \alpha \end{bmatrix} \bar{M} \begin{bmatrix} 1 \\ \alpha \end{bmatrix} \end{aligned}$$

This condition can be stated linearly in terms of the coefficients of M, \bar{M} and those of $P(\cdot)$. Denoting the i, j^{th} entrance of M by M_{ij} , $i, j \in [2]$, the list

of constraints is given by

$$\begin{aligned} \frac{2\mu}{2L + \mu} \bar{M}_{11} &= -a_0, & M_{11} - \bar{M}_{11} + \frac{4\mu}{2L + \mu} \bar{M}_{12} &= -a_1 \\ 2M_{12} - 2\bar{M}_{12} + \frac{2\mu}{2L + \mu} \bar{M}_{22} &= -a_2, & M_{22} - \bar{M}_{22} &= -a_3. \end{aligned}$$

□

5.B Proof of Proposition 5.7

Proof. Condition (5.11) can be equivalently expressed as

$$0 \geq (1 - \lambda)A_{ijm} + \lambda A_{ikm} + \lambda(1 - \lambda)B_{jkm} := P(\lambda), \quad \forall m \in [n], \quad \forall \lambda \in [0, 1].$$

By Markov-Lukasc theorem [Sze39, Theorem 1.21],

$$\begin{aligned} P(\lambda) &\leq 0, \quad \forall \lambda \in [0, 1] \\ \Leftrightarrow \exists M \in \mathbb{S}_+^2, K \geq 0 \text{ s.t. } P(\lambda) &= - \begin{bmatrix} 1 & \lambda \end{bmatrix} M \begin{bmatrix} 1 \\ \lambda \end{bmatrix} - \lambda(1 - \lambda)K. \end{aligned}$$

This condition can be stated linearly in terms of the coefficients of M , of K and of the coefficients of $P(\cdot)$. Denoting the i, j^{th} entrance of M by M_{ij} , $i, j \in [2]$, the list of constraints is given by

$$M_{11} = -A_{ijm}, \quad 2M_{12} + K = A_{ijm} - A_{ikm} - B_{jkm}, \quad M_{22} - K = B_{jkm}.$$

□

5.C Proof of Proposition 5.9

Proof. Take the first condition in (5.15), and set $\lambda_k = \gamma^2$, $\mu_i = \theta^2$ and $\mu_j = \alpha^2$. The condition becomes

$$\begin{aligned} 0 &\geq \gamma^2(\alpha^2 + \theta^2)(A_{ik}\theta^2 + A_{jk}\alpha^2) - \gamma^2\theta^2\alpha^2 A_{ij} \\ &\quad + \gamma^4(B_{ik}\theta^2 + B_{jk}\alpha^2) + B_{ij}\theta^2\alpha^2(\theta^2 + \alpha^2) = P(\alpha, \theta, \gamma). \end{aligned}$$

$P(\alpha, \theta, \gamma)$ is an homogeneous polynomial of 3 variables and degree 6, which must be negative for all $\theta, \alpha, \gamma \in \mathbb{R}$. This negativity condition can be relaxed by requiring P to be Sum of Squares, that is, by requiring the existence of some $M \preceq 0$ such that

$$P(\alpha, \theta, \gamma) = V(\alpha, \theta, \gamma)^\top M V(\alpha, \theta, \gamma),$$

where $V(\alpha, \theta, \gamma) = (\alpha^2\theta \ \alpha^2\gamma \ \alpha\theta^2 \ \alpha\gamma^2 \ \alpha\theta\gamma \ \theta^2\gamma \ \theta\gamma^2)^\top$. This condition can be stated linearly in terms of the coefficients of M and those of $P(\cdot)$. Denoting the i, j^{th} entrance of M by $M_{i,j}$, $i, j \in [7]$, the list of constraints is given by

$$M_{1,1} = M_{3,3} = B_{ij}$$

$$M_{2,2} = A_{jk}$$

$$M_{4,4} = B_{jk}$$

$$M_{6,6} = A_{ik}$$

$$M_{7,7} = B_{ik}$$

$$M_{5,5} + 2M_{1,7} + 2M_{2,6} + 2M_{3,4} = -A_{ij} - 2\mu B_{ij} + A_{ik} + A_{jk}$$

$$M_{1,2} = M_{1,3} = M_{2,4} = M_{3,6} = M_{4,7} = M_{6,7} = 0$$

$$M_{1,5} + M_{2,3} = M_{1,4} + M_{2,5} = M_{1,6} + M_{3,5} = 0$$

$$M_{2,7} + M_{4,5} = M_{3,7} + M_{5,6} = M_{4,6} + M_{5,7} = 0.$$

Translating the second condition in (5.15) follows exactly the same argument, with $A_{(\cdot)}$ and $B_{(\cdot)}$ inverted. \square

PART II
Interpolation conditions for
second-order univariate
optimization

6

Generic method for univariate interpolation

This chapter results from a collaboration with Nizar Bousselmi.

Suppose interpolation conditions are known for a class \mathcal{F} of *univariate* functions. This chapter proposes a technique, summarized in Theorem 6.1, to derive interpolation conditions for the associated univariate class

$$\int \mathcal{F} := \{f : \mathbb{R} \rightarrow \bar{\mathbb{R}} \mid f' \in \mathcal{F}\}, \quad (6.1)$$

consisting of functions whose derivative belongs to \mathcal{F} . This approach thus lifts interpolation conditions from \mathcal{F} to the higher-order class $\int \mathcal{F}$. It will be relied upon in Chapter 7 to derive interpolation conditions for univariate second-order function classes, enabling tight analysis of univariate second-order optimization methods.

This chapter is organized as follows. Section 6.1 generalizes the definition of an interpolation condition to univariate high-order function classes. Section 6.2 then introduces the method on the simple example of M -smooth functions, and Sections 6.3, 6.4, and 6.5 formalize the method, summarized in Section 6.6.

Notation We denote by \mathcal{C}^m the class of univariate functions $f : \mathbb{R} \rightarrow \mathbb{R}$ that are at least m times everywhere differentiable, with derivative up to order m everywhere continuous. In addition, in Section 7.3, we will consider function classes whose elements f may take infinite values, but such that their inverse $\frac{1}{f(x)} \in \mathcal{C}^{(0)}$. To account for such function classes, we introduce the class $\bar{\mathcal{C}}^0$ of functions $f : \mathbb{R} \rightarrow \bar{\mathbb{R}}$ that are everywhere continuous, i.e., $\forall c \in \mathbb{R}, \lim_{x \rightarrow c} f(x) = f(c)$, as in

$$\begin{cases} \forall \epsilon > 0, \exists \delta > 0, \forall x \in \mathbb{R} : |x - c| < \delta \Rightarrow |f(x) - f(c)| < \epsilon & f(c) \text{ finite,} \\ \forall N > 0, \exists \delta > 0, \forall x \in \mathbb{R} : |x - c| < \delta \Rightarrow f(x) > N & f(c) \text{ infinite.} \end{cases}$$

Similarly, we introduce classes $\bar{\mathcal{C}}^{(m)}$ to account for possibly infinite-valued functions whose inverse $\frac{1}{f(x)} \in \mathcal{C}^{(m)}$. Let $f \in \bar{\mathcal{C}}^0$, and suppose f is at least m times differentiable on $\text{dom } f$. We denote by $f^{(k)}$ the k^{th} derivative of f (where $f^{(0)} = f$), defined $\forall c \in \mathbb{R}$ and for $k \geq 1$ as

$$f^{(k)}(c) = \lim_{x \rightarrow c} \frac{f^{(k-1)}(x) - f^{(k-1)}(c)}{x - c},$$

i.e., as the classical order k derivative on the effective domain, and $+\infty$ otherwise. We say $f \in \bar{\mathcal{C}}^m$ if (i) $f^{(k)}$ exists, and (ii) $f^{(k)} \in \bar{\mathcal{C}}^0, \forall k \leq m$.

6.1 Order m interpolation conditions

We seek for interpolation constraints for classes of univariate functions $\mathcal{F} \subseteq \bar{\mathcal{C}}^m$, hence generalize Definition 2.6 to the interpolation up to order m derivative of datasets by functions in \mathcal{F} .

Definition 6.1 (Interpolation with function values). Given $\mathcal{F} \subseteq \bar{\mathcal{C}}^m$, a set $S = \{(x_i, f_i^0, f_i^1, \dots, f_i^m)\}_{i \in [N]} \in (\mathbb{R} \times \dots \times \mathbb{R})^N$ is \mathcal{F} -interpolable with function values if and only if

$$\exists f \in \mathcal{F} : f^{(k)}(x_i) = f_i^k, \forall k \in [m], \quad \forall i \in [N]. \tag{6.2}$$

Definition 6.1 can be straightforwardly extended to interpolation without function values, that is, interpolation of a dataset S by a function in \mathcal{F} , except for its function values f_i^0 .

Definition 6.2 (Interpolation without function values). Given $\mathcal{F} \subseteq \bar{\mathcal{C}}^m$, a set $S = \{(x_i, f_i^0, f_i^1, \dots, f_i^m)\}_{i \in [N]} \in (\mathbb{R} \times \dots \times \mathbb{R})^N$ is \mathcal{F} -interpolable without function values if and only if

$$\exists f \in \mathcal{F} : f^{(k)}(x_i) = f_i^k, \quad \forall k \in [m], \forall i \in [N]. \quad (6.3)$$

The quantities f_i^0 are thus ignored in interpolation without function values and are not necessarily equal to function values $f^{(0)}(x_i)$.

Remark 6.1. In the multivariate case, interpolation conditions without function values are often computationally more expensive to impose than interpolation conditions with function values, see, e.g., [Tay17, Section 3.4]. By contrast, in the univariate case, these interpolation conditions are often more concise than interpolation conditions with function values. Hence, we resort to this formulation whenever possible, for instance when analyzing methods that do not involve function values in their settings.

Whenever $m = 1, 2$, we denote $f_i^0 := f_i, f_i^1 := g_i, f_i^2 := h_i, f^{(1)}(x) := f'(x)$ and $f^{(2)}(x) := f''(x)$.

6.2 Overview of the method

This section illustrates the main steps of the proposed method on a simple example, depicted in Figure 6.1. Consider the class \mathcal{F}_M of M -Lipschitz functions, whose interpolation conditions correspond to a discretization of their classical characterization, i.e., a dataset $S = \{(x_i, f_i)\}_{i \in [N]}$ is \mathcal{F}_M -interpolable if and only if

$$|f_i - f_j| \leq M|x_i - x_j|, \quad \forall i, j \in [N] \text{ [Kir34, Val45]}.$$

Building on these conditions, we aim to recover interpolation conditions for the class $\int \mathcal{F}_M$ of M -smooth functions, that is conditions ensuring $\int \mathcal{F}_M$ -interpolability of a dataset $S = \{(x_i, f_i, g_i)\}_{i \in [N]}$.

First, we address interpolation conditions without function values for $\int \mathcal{F}_M$. These conditions coincide with the interpolation conditions with function values for \mathcal{F}_M (Lemma 6.1). Specifically, S is $\int \mathcal{F}_M$ -interpolable without function values if and only if

$$|g_i - g_j| \leq M|x_i - x_j|, \quad \forall i, j \in [N]. \quad (6.4)$$

Indeed, then, there exists a function in \mathcal{F}_M interpolating $\{(x_i, g_i)\}$, and this function is, by the fundamental theorem of calculus, the derivative of some function in $\int \mathcal{F}_M$.

Second, we exploit these interpolation conditions without function values to derive interpolation conditions *with* function values for $\int \mathcal{F}_M$. Since continuous piecewise smooth functions are globally smooth, the problem reduces to interpolation on a single interval $[x_i, x_j]$ (see Lemma 6.5). Suppose, on this interval, S satisfies (6.4). Consider all functions in \mathcal{F}_M that interpolate $\{(x_i, g_i), (x_j, g_j)\}$, and in particular the *extremal* interpolants g_{\min} and g_{\max} (Definition 6.3), see Figure 6.1 for an illustration.




Fig. 6.1 Illustration of the method to obtain interpolation conditions ensuring $S = \{(x_i, f_i, g_i)\}_{i=1,2}$ to be \mathcal{F}_M -interpolable: extremal M -Lipschitz interpolants g_{\min} (top figure, dark blue curve) and g_{\max} (top figure, pink curve) of $\{(x_i, g_i)\}_{i=1,2}$ (see Proposition 7.3 for their derivation). When integrated (below, dark blue and pink curves), these extremal interpolants define an interval including all possible smooth functions interpolating S , except f_2 . The interval $[f_1 + \int_{x_1}^{x_2} g_{\min}(x)dx, f_1 + \int_{x_1}^{x_2} g_{\max}(x)dx]$ (bottom figure, light blue line) is exactly the interval of admissible values for f_2 . Whenever f_2 belongs to this interval, there exists a convex combination of the integrals of g_{\min} and g_{\max} (light blue dotted curves) which is in $\int \mathcal{F}_M$ and interpolates S .

Given f_i , integration of these extremal interpolants over $[x_i, x_j]$ yields

the interval

$$\left[f_i + \int_{x_i}^{x_j} g_{\min}(x) \, dx, \quad f_i + \int_{x_i}^{x_j} g_{\max}(x) \, dx \right]$$

of admissible values for f_j (cf. Lemma 6.4). Indeed, if f_j lies within this interval, there exists a convex combination $g = \lambda g_{\min} + (1 - \lambda) g_{\max}$ such that the function $f_i + \int_{x_i}^x g(t) \, dt$ (i) is M -smooth by convexity of the smoothness property, i.e., any convex combination of smooth continuous functions is itself smooth, and (ii) interpolates S . Conversely, if f_j lies outside this interval, no function in \mathcal{F}_M interpolates S . Hence, imposing the interpolation condition

$$\int_{x_i}^{x_j} g_{\min}(x) \, dx \leq f_j - f_i \leq \int_{x_i}^{x_j} g_{\max}(x) \, dx, \quad (6.5)$$

in addition to the conditions without function values, yields interpolation conditions with function values for $\int \mathcal{F}_M$, see Corollary 7.1.

In this toy example, to recover interpolation conditions for $\int \mathcal{F}_M$ from those of \mathcal{F}_M , we relied on several key properties of both \mathcal{F}_M and $\int \mathcal{F}_M$. More specifically, the proposed method applies to classes \mathcal{F} satisfying

1. $\mathcal{F} \subseteq \mathcal{C}^m$ is a class of univariate continuous functions (obtention of interpolation condition without function values for $\int \mathcal{F}$),
2. \mathcal{F} admits well-defined extremal interpolants, ensured by Assumption 6.1 (definition of an interval of possible values for f_2),
3. $\int \mathcal{F}$ is convex, or satisfies the weaker notion introduced in Assumption 6.2 (the interval whose extremes are the integrated extremal extrapolants of \mathcal{F} indeed contains all possible values for f_2),
4. Any function in \mathcal{C}^m that belongs, piecewise, to $\int \mathcal{F}$, is in $\int \mathcal{F}$, ensured by Assumption 6.4 (sufficiency of piecewise $\int \mathcal{F}$ -interpolability for global $\int \mathcal{F}$ -interpolability).

We now detail the steps of the method and underlying assumptions. For illustrative applications of the technique to a broad range of function classes with second-order properties, including, e.g., functions with Lipschitz continuous Hessian, see Chapter 7.3.

6.3 Interpolation without function values

In the univariate case and for continuous function classes, interpolation without function values of $\int \mathcal{F}$ amounts exactly to interpolation with function values of \mathcal{F} ,

Lemma 6.1 (Interpolation without function values). *Consider $\mathcal{F} \subseteq \bar{\mathcal{C}}^m$ and $\int \mathcal{F} \subseteq \bar{\mathcal{C}}^{m+1}$, as defined in (6.1). A set $S = \{(x_i, f_i^0, f_i^1, \dots, f_i^m)\}_{i \in [N]} \in (\mathbb{R} \times \dots \times \mathbb{R})^N$ is $\int \mathcal{F}$ -interpolable (without function values) if and only if $\tilde{S} := \{(x_i, f_i^1, \dots, f_i^m)\}_{i \in [N]}$ is \mathcal{F} -interpolable (with function values).*

Proof. By Definition 6.2, S is $\int \mathcal{F}$ -interpolable without function values if

$$\begin{aligned} & \exists f \in \int \mathcal{F} && : f^{(k)}(x_i) = f_i^k, \quad \forall k \in [m], \forall i \in [N] \\ \Leftrightarrow & \exists f : f' = g \in \mathcal{F}, f^{(k)}(x_i) = f_i^k, \quad \forall k \in [m], \forall i \in [N] \\ (6.1) & && \\ \Leftrightarrow & \exists g \in \mathcal{F} && : g^{(k)}(x_i) = f_i^{k+1}, \quad \forall k = 0, \dots, m-1, \forall i \in [N] \\ \Leftrightarrow & \tilde{S} \text{ is } \mathcal{F}\text{-interpolable with function values,} \\ \text{Def. 6.1} & && \end{aligned}$$

where the second equality follows from the fundamental theorem of calculus, see, e.g., [Rud76, Chapter 6] applied on the effective domain of g . \square

6.4 Interpolation with function values (2 points)

We now show how to go from interpolation without function values to interpolation with function values. We take the simplest situation possible and restrict our attention to the interpolation of a single pair of data points $x_1 < x_2$.

Consider a set $S = \{(x_i, f_i^0, f_i^1, \dots, f_i^{m+1})\}_{i=1,2}$ which is $\int \mathcal{F}$ -interpolable without function values and the associated set $\tilde{S} = \{(x_i, f_i^1, \dots, f_i^{m+1})\}_{i=1,2}$. By Lemma 6.1, \tilde{S} is \mathcal{F} -interpolable, hence there exists at least one function $g \in \mathcal{F}$ interpolating \tilde{S} . For each of such g , we can construct an associated function $f \in \int \mathcal{F}$, interpolating S with the exception of $f(x_2)$, which could differ from f_2^0 .

Lemma 6.2. Consider $\mathcal{F} \subseteq \bar{\mathcal{C}}^m$, $\int \mathcal{F} \in \bar{\mathcal{C}}^{m+1}$ as defined in (6.1), and a set $S = \{(x_i, f_i^0, f_i^1, \dots, f_i^{m+1})\}_{i=1,2} \in (\mathbb{R} \times \dots \times \mathbb{R})^2$, $\int \mathcal{F}$ -interpolable without function values. Let $\tilde{S} = \{(x_i, f_i^1, \dots, f_i^{m+1})\}_{i=1,2}$, and $g \in \mathcal{F}$ be a function interpolating \tilde{S} . Then,

$$f(x) := f_1^0 - \int_{-\infty}^{x_1} g(z) dz + \int_{-\infty}^x g(z) dz, \quad (6.6)$$

belongs to $\int \mathcal{F}$ and interpolates S except $f(x_2)$ which could differ from f_2^0 .

Proof. Since $f'(x) = g(x)$, $f \in \int \mathcal{F}$ by (6.1). In addition,

$$f^{(k)}(x_i) = g^{(k-1)}(x_i) = f_i^k, \quad i = 1, 2, \quad k \in [m+1], \quad \text{and} \quad f(x_1) = f_1^0,$$

hence f interpolates S except possibly for $f(x_2)$. □

Consider, among these functions $g \in \mathcal{F}$ interpolating S , the lowest and highest functions, called the *extremal interpolants* of S , see Figure 6.1 for an illustration. As such, extremal interpolants are not well-defined and might not exist or be non-unique. We thus introduce the notion of *lower and upper interpolating envelopes*, that is pointwise minimum and maximum of all functions in \mathcal{F} interpolating S .

Definition 6.3 (Extremal interpolating envelopes). Consider $\mathcal{F} \subseteq \bar{\mathcal{C}}^m$ and an \mathcal{F} -interpolable set $S = \{(x_i, g_i^0, g_i^1, \dots, g_i^m)\}_{i=1,2} \in (\mathbb{R} \times \dots \times \mathbb{R})^2$, where $x_1 < x_2$. The *extremal interpolating envelopes* g_{\min} and g_{\max} of S are defined as, $\forall x \in \mathbb{R}$,

$$\begin{aligned} g_{\min} \text{ (resp. } \max) (x) &= \inf_{g: \mathbb{R} \rightarrow \mathbb{R}} \text{ (resp. } \sup) \quad g(x) \\ \text{s.t. } &g \in \mathcal{F} \text{ and } g^{(k)}(x_i) = g_i^k, \quad i = 1, 2, \quad k \in [m]. \end{aligned} \quad (6.7)$$

Extremal interpolating envelopes always exist, but do not always belong to \mathcal{F} . To ensure it is the case, we require \mathcal{F} to be *extremally interpolable*.

Assumption 6.1 (Extremally interpolable function class). We say a function class $\mathcal{F} \subseteq \bar{\mathcal{C}}^m$ of univariate functions is *extremally interpolable* if any \mathcal{F} -interpolable set $S = \{(x_i, g_i^0, g_i^1, \dots, g_i^m)\}_{i=1,2} \in (\mathbb{R} \times \dots \times \mathbb{R})^2$ satisfies $g_{\min}, g_{\max} \in \mathcal{F}$ where g_{\min} and g_{\max} are defined in (6.7).

Whenever \mathcal{F} is extremally interpolable, we refer to extremal interpolating envelopes of S as extremal interpolants of S . Showing extremal interpolability of \mathcal{F} often amounts to computing these extremal interpolants.

Suppose now that such extremal interpolants exist, and that the associated functions $f_{\min}, f_{\max} \in \int \mathcal{F}$ as defined in (6.6) evaluated at x_2 are respectively smaller and bigger than f_2^0 . Then, there exists a function $f \in \int \mathcal{F}$ interpolating S (including $f(x_2) = f_2^0$), provided $\int \mathcal{F}$ is *extremally completable*.

Assumption 6.2 (Extremally completable function class). Let $\mathcal{F} \subseteq \bar{\mathcal{C}}^m$ be extremally interpolable (Assumption 6.1), and consider $\int \mathcal{F} \in \bar{\mathcal{F}}^{m+1}$. Let $S = \{(x_i, f_i^0, f_i^1, \dots, f_i^{m+1})\}_{i=1,2} \in (\mathbb{R} \times \dots \times \mathbb{R})^2$, where $x_1 \leq x_2$, be $\int \mathcal{F}$ -interpolable without function values, $\tilde{S} = \{(x_i, f_i^1, \dots, f_i^{m+1})\}_{i=1,2}$, and $g_{\min}, g_{\max} \in \mathcal{F}$ be the extremal interpolants of \tilde{S} , as defined in (6.7). Let, in addition,

$$f_1^0 + \int_{x_1}^{x_2} g_{\min}(z) dz \leq f_2^0 \leq f_1^0 + \int_{x_1}^{x_2} g_{\max}(z) dz.$$

We then say $\int \mathcal{F}$ is *extremally completable* if

$$\exists f \in \int \mathcal{F} : f^{(k)}(x_i) = f_i^k, \quad \forall i \in \{1, 2\}, \quad \forall k \in [m + 1].$$

Instances of function classes $\int \mathcal{F}$ that are extremally completable include any function class satisfying the maybe more natural notion of *convexity*, including quasi-self-concordant functions and functions with Lipschitz Hessian, see Proposition 7.2.

Assumption 6.3 (Convex function class). We say a class \mathcal{F} of univariate functions is *convex* if, given any $f_a, f_b \in \mathcal{F}$ and any $\lambda \in [0, 1]$,

$$\lambda f_a + (1 - \lambda) f_b \in \mathcal{F}.$$

Convex function classes are extremally completable.

Lemma 6.3. Let $\mathcal{F} \subseteq \bar{\mathcal{C}}^m$ be an extremally interpolable (Assumption 6.1) class of univariate functions. If $\int \mathcal{F}$ is convex (Assumption 6.3), then $\int \mathcal{F}$ is extremally completable (Assumption 6.2).

Proof. With the same notation as in Assumption 6.2, let

$$f_{\min/\max}(x) := f_1^0 - \int_{-\infty}^{x_1} g_{\min/\max}(z) dz + \int_{-\infty}^x g_{\min/\max}(z) dz.$$

By Lemma 6.2, $f_{\min}(x)$ and $f_{\max}(x)$ belong to $\int \mathcal{F}$. In addition, the two functions interpolate S with the exception of $f_2^0 = \lambda f_{\text{lo}}(x_2) + (1 - \lambda) f_{\text{hi}}(x_2)$

for some $\lambda \in [0,1]$. Hence $f(x) := \lambda f_{\text{lo}}(x) + (1 - \lambda)f_{\text{hi}}(x)$ interpolates S (including $f(x_2) = f_2$), and by convexity of $\int \mathcal{F}$, $f \in \int \mathcal{F}$. \square

Convexity on its own proves to be too restrictive, since not satisfied by classes of interest, e.g., self-concordant functions. This justifies the introduction of the weaker notion of extremal completability, satisfied by self-concordant functions.

For classes \mathcal{F} that are extremally interpolable, and such that $\int \mathcal{F}$ is extremally completable, we obtain interpolation conditions with function values for $\int \mathcal{F}$, on a single interval, by enforcing f_2^0 to belong to the interval $[f_1^0 + \int_{x_1}^{x_2} g_{\min}(z)dz, f_1^0 + \int_{x_1}^{x_2} g_{\max}(z)dz]$ defined by the integrals of the extremal interpolants of S .

Lemma 6.4 (Interpolation conditions for a single interval). *Let $\mathcal{F} \subseteq \bar{\mathcal{C}}^m$ be an extremally interpolable (Assumption 6.1) class of univariate functions and let $\int \mathcal{F} \subseteq \bar{\mathcal{C}}^m$ (defined in (6.1)) be extremally completable (Assumption 6.2). A set $S = \{(x_i, f_i^0, f_i^1, \dots, f_i^{m+1})\}_{i=1,2} \in (\mathbb{R} \times \dots \times \mathbb{R})^2$, where $x_1 < x_2$ is $\int \mathcal{F}$ -interpolable with function values if and only if S is $\int \mathcal{F}$ -interpolable without function values, and f_1^0, f_2^0 satisfy*

$$\int_{x_1}^{x_2} g_{\min}(x)dx \leq f_2^0 - f_1^0 \leq \int_{x_1}^{x_2} g_{\max}(x)dx, \tag{6.8}$$

where g_{\min} and g_{\max} are defined as in (6.7).

Proof. Sufficiency follows the definition of extremal completability. To prove necessity, suppose $S = \{(x_i, f_i^0, f_i^1, \dots, f_i^{m+1})\}_{i=1,2}$ is $\int \mathcal{F}$ -interpolable (with function values). Then, by Lemma 6.1, S is \mathcal{F} -interpolable (without function values). In addition, by Definition 6.3 of extremal interpolants, any function $g : \mathbb{R} \rightarrow \mathbb{R}$ such that (i) $\exists f \in \int \mathcal{F} : g(x) = f'(x), \forall x \in \mathbb{R}$ and (ii) g interpolates S , satisfies

$$g_{\min}(x) \leq g(x) \leq g_{\max}(x), \forall x \in [x_1, x_2].$$

Hence, (6.8) is necessarily satisfied. \square

To obtain full interpolation conditions for $\int \mathcal{F}$, it remains to consider the case of an arbitrary number of points to interpolate.

6.5 Interpolation with function values (all points)

In the univariate case, interpolation of an arbitrary number of points is equivalent to that of a single pair, provided the class $\int \mathcal{F}$ is *order $m + 1$ connectable*, in the sense that the juxtaposition of different functions in $\int \mathcal{F}$, where the functions coincide up to order $m + 1$ on the boundaries of some intervals, is itself a function in $\int \mathcal{F}$.

Assumption 6.4 (Order m connectable function class). A class $\mathcal{F} \subseteq \bar{\mathcal{C}}^m$ of univariate functions is *order m connectable* if, given any $x_0 \leq \dots \leq x_{K+1} \in \mathbb{R}$, and any K functions $f_j \in \mathcal{F}$ such that

$$f_j^{(l)}(x_{j+1}) = f_{j+1}^{(l)}(x_{j+1}), \quad \forall l \in [m], \quad j \in [K-1],$$

then the piecewise function

$$f: \mathbb{R} \rightarrow \bar{\mathbb{R}}: f(x) := f_j(x) \quad \forall x \in \mathcal{I}_j, \quad \forall j \in [K] \text{ belongs to } \mathcal{F},$$

where

$$\mathcal{I}_j = \begin{cases}]-\infty, x_1] & j = 1 \\ [x_j, x_{j+1}] & j = 2, \dots, K-1 \\ [x_K, \infty[& j = K. \end{cases}$$

Remark 6.2. Linear functions are order 1 connectable, since the juxtaposition of several linear functions whose slope is identical is linear, but not order 0 connectable, since the juxtaposition of linear functions with different slopes is not linear. Hence, an appropriate choice of order is essential when determining whether a function class is connectable.

Order $m + 1$ connectivity of $\int \mathcal{F}$ allows extending Lemma 6.4 to any arbitrary number of points.

Lemma 6.5 (From a single interval to all points). *Let $\mathcal{F} \subseteq \bar{\mathcal{C}}^m$ be an order m connectable class of univariate functions (Assumption 6.4). A set $S = \{(x_i, f_i^0, f_i^1, \dots, f_i^m)\}_{i \in [N]} \in (\mathbb{R} \times \dots \times \mathbb{R})^N$, with ordered points $x_0 \leq x_1 \leq \dots \leq x_N$, is \mathcal{F} -interpolable with function values if and only if,*

$$\forall i \in [N-1]: \left\{ (x_i, f_i^0, f_i^1, \dots, f_i^m), (x_{i+1}, f_{i+1}^0, f_{i+1}^1, \dots, f_{i+1}^m) \right\} \text{ is } \mathcal{F}\text{-interpolable.} \quad (6.9)$$

Proof of Lemma 6.5. (Necessity) Suppose S is \mathcal{F} -interpolable, and let $f \in \mathcal{F}$ be a function interpolating S , and in particular all pairs (i, j) . For each pair (i, j) , the restriction of f to the interval $[x_i, x_j]$ is a function in \mathcal{F} interpolating $\{(x_i, f_i^0, f_i^1, \dots, f_i^m), (x_j, f_j^0, f_j^1, \dots, f_j^m)\}$, and (6.9) is satisfied.

(Sufficiency) Suppose (6.9) is satisfied, and let $f^{(i)} \in \mathcal{F}$ be a function interpolating the pair $(i, i+1) \forall i \in [N-1]$. Define

$$\mathcal{I}_i = \begin{cases}]-\infty, x_2] & i = 1 \\ [x_i, x_{i+1}] & i = 2, \dots, N-2 \\ [x_{N-1}, \infty[& i = N-1. \end{cases} \quad (6.10)$$

Let $f : \mathbb{R} \rightarrow \bar{\mathbb{R}}$ be defined by:

$$f(x) := f^{(i)}(x) \quad \forall x \in \mathcal{I}_i, \forall i \in [N-1]. \quad (6.11)$$

By construction, f interpolates $S = \{(x_i, f_i^0, f_i^1, \dots, f_i^m)\}_{i \in [N]}$. Further, by order m piecewise invariance of \mathcal{F} , $f \in \mathcal{F}$, which completes the proof. \square

6.6 Main Theorem

We now present our main Theorem, lifting interpolation conditions from \mathcal{F} to $\int \mathcal{F}$.

Theorem 6.1 (Interpolation conditions for $\int \mathcal{F}$)

Let $\mathcal{F} \subseteq \bar{\mathcal{C}}^m$ be an extremally interpolable (Assumption 6.1) class of univariate functions, and let $\int \mathcal{F} \subseteq \bar{\mathcal{C}}^{m+1}$ (defined in (6.1)) be extremally completable (Assumption 6.2) and order $m+1$ connectable (Assumption 6.4). A set $S = \{(x_i, f_i^0, f_i^1, \dots, f_i^m)\}_{i \in [N]} \in (\mathbb{R} \times \dots \times \mathbb{R})^N$, where $x_0 \leq x_1 \leq \dots \leq x_N$ is $\int \mathcal{F}$ -interpolable if and only if S is \mathcal{F} -interpolable without function values, and $\forall i \in [N]$,

$$\int_{x_i}^{x_{i+1}} g_{\min}(x) dx \leq f_{i+1} - f_i \leq \int_{x_i}^{x_{i+1}} g_{\max}(x) dx, \quad (6.12)$$

where g_{\min}, g_{\max} are defined as in (6.7).

Proof. It suffices to combine Lemmas 6.4 and 6.5.

□

Remark 6.3. Theorem 6.1 requires imposing conditions on ordered pairs only. However, imposing conditions on all pairs yields, by Lemma 6.5, equivalent interpolation conditions, that are often cleaner and easier to handle.

7

Computer-aided analysis of second-order univariate optimization

This chapter results from a collaboration with Nizar Bousselmi.

As stated in Chapter 3, computer-aided techniques for tight performance analysis of optimization methods are to this day mostly restricted to the analysis of first-order optimization. In this setting, obtaining performance guarantees often amounts to solving SDPs, i.e., the resulting PEP are convex; and interpolation conditions have been derived for a wide range of first-order function classes, ensuring the tightness of these guarantees.

By contrast, we seek to analyze second-order optimization methods, which, to solve unconstrained minimization problems

$$\min_{x \in \mathbb{R}^d} f(x),$$

rely on evaluations of the Hessian of the objective function at each iteration, e.g., Newton's method, whose iteration is given by

$$x_{k+1} = x_k - \nabla^2 f(x_k)^{-1} \nabla f(x_k). \quad (\text{NM})$$

Such methods are designed to minimize functions belonging to second-order function classes, that is, classes whose characterization involves evaluations of second-order quantities, e.g, the class of twice-differentiable functions $f : \mathbb{R}^d \rightarrow \mathbb{R}^d$ whose Hessian is M -Lipschitz continuous, i.e,

$$\|\nabla^2 f(x) - \nabla^2 f(y)\| \leq M\|x - y\|, \quad \forall x, y \in \mathbb{R}^d. \quad (7.1)$$

This chapter takes a first step at leveraging the PEP framework to the analysis of second-order optimization, by (i) obtaining interpolation conditions for *univariate* second-order function classes, (Section 7.3), and (ii) solving the associated PEPs, (Section 7.4). Before diving into these main contributions, we quickly motivate the analysis of (univariate) second-order methods, and review computer-aided techniques in this setting.

7.1 Codes.

Codes for reproducing and verifying technical parts of this work can be found at

<https://github.com/NizarBousselmi/Second-Order-Univariate-PEP>.

7.2 Motivation

7.2.1 Second-order optimization

Second-order methods such as Newton's method have attracted a lot of attention due to their local quadratic convergence, and importance for interior-point methods [NN94], widely used to solve constrained optimization problems, see, e.g., [CGT00] for a survey.

Since Newton's method may globally diverge, and fails on degenerate Hessians, several globally convergent variants of Newton's method have been designed that handle these issues. These mostly include (i) *damped* Newton's methods, i.e.,

$$x_{k+1} = x_k - \alpha_k \nabla^2 f(x_k)^{-1} \nabla f(x_k), \quad (\text{DNM})$$

and (ii) *regularized* Newton's methods [Mar63, Lev44], i.e.,

$$x_{k+1} = x_k - \alpha_k (\nabla^2 f(x_k) + \lambda_k \text{Id}_d)^{-1} \nabla f(x_k), \quad (\text{RNM})$$

where λ_k may depend on x_k , x_{k+1} , or $\nabla f(x_k)$, and Id_d is the identity matrix in \mathbb{S}^d .

These methods have been designed for and analyzed on several function classes, including those of functions with Lipschitz Hessian, denoted \mathcal{H}_M ; of convex functions therein, denoted $\mathcal{H}_{M,0}$; of self-concordant functions [NN94], denoted $\mathcal{S}_{M,+}$, that is, convex functions satisfying, in the univariate case,

$$|f'''(x)| \leq 2M(f''(x))^{3/2}, \quad \forall x \in \mathbb{R}^d, \quad (7.2)$$

and including e.g., barrier functions, that form the basis of interior-point methods; and of quasi-self-concordant functions [Bac10], denoted $\mathcal{T}_{M,+}$, i.e., convex functions satisfying, still in the univariate case,

$$|f'''(x)| \leq Mf''(x), \quad \forall x, h \in \mathbb{R}^d, \quad (7.3)$$

and including logistic regressions.

Examples of globally convergent second-order methods on these classes include,

- Cubic Regularized Newton method (CNM), i.e., $\lambda_k \propto \|x_{k+1} - x_k\|$, on \mathcal{H}_M [NP06], achieving $\mathcal{O}(1/k^{\frac{2}{3}})$ convergence; accelerated CNM on $\mathcal{H}_{M,0}$ [Nes08] showing $\mathcal{O}(1/k^4)$ convergence; adaptive CNM on functions with milder assumptions than Lipschitz continuity on the Hessian [CGT11];
- Gradient regularized Newton methods (GNM1)- $\lambda_k \propto \|\nabla f(x_k)\|$, on convex functions [Pol09], achieving $\mathcal{O}(1/k^{\frac{1}{4}})$ convergence; and $\mathcal{T}_{M,+}$ [Doi23], achieving linear convergence;
- Gradient regularized Newton methods (GNM2)- $\lambda_k \propto \|\nabla f(x_k)\|^2$, on \mathcal{H}_M and $\mathcal{H}_{M,0}$ [DN24, Mis23], achieving $\mathcal{O}(1/k^2)$ convergence;
- Damped Newton method on $\mathcal{S}_{M,+}$ [NN94, Hil21, IH24], achieving linear convergence; and on semi-self-concordant functions [HKP⁺22], achieving $\mathcal{O}(1/k^2)$ convergence; and
- Universal methods [DMN24, DN21].

In general, the potential tightness of existing convergence rates is an open question. In practice, the methods often behave significantly better than predicted by the theory, either because worst-case instances are very unlikely to appear in practical problems, or because the current convergence rates are too pessimistic and can be improved, e.g. by relying on interpolation conditions. In addition, a fair comparison between the different variants of Newton's method is difficult to perform, in the sense that the performance measures and hypotheses vary for each analysis.

7.2.2 Computer-aided second-order optimization

The theory underlying the PEP framework can be straightforwardly extended to second-order optimization. For instance, the worst-case analysis of N iterations of (NM) on \mathcal{H}_M , where the performance measure is chosen to be $\frac{\|x_N - x_\star\|}{\|x_0 - x_\star\|}$, can be expressed as

$$\begin{aligned}
 & \max_{S = \{(x_k, f_k, g_k, h_k)\}_{k=0,1,\dots,N,\star}} \|x_N - x_\star\| & (7.4) \\
 \text{s.t. } & x_{k+1} = x_k - h_k^{-1} g_k, \quad k = 0, 1, \dots, N-1 \\
 & \|x_0 - x_\star\| \leq R \\
 & g_\star = 0, h_\star \succcurlyeq 0 \\
 & \exists f \in \mathcal{H}_M : f(x_k) = f_k, \nabla f(x_k) = g_k, \nabla^2 f(x_k) = h_k, \quad k = 0, 1, \dots, N, \star.
 \end{aligned}$$

As compared to the analysis of first-order methods, (7.4) presents two difficulties. First, no interpolation condition is yet known for \mathcal{H}_M , nor for any second-order function class, see Chapter 2, preventing any exact algebraic representation of the last constraint in \mathcal{H}_M . Second, by contrast to the majority of first-order problems, the constraints in (7.4) are not Gram representable, hence (7.4) is a non-convex PEP with, in the multivariate case, matrix variables.

This second difficulty can be tackled by leveraging, e.g., techniques to solve non-convex PEPs yielded by first-order problems, e.g., in [DGFST24, DGVPR24, RTBG20].

Regarding interpolation conditions, to the best of our knowledge, no progress has been achieved so far. The only attempts at computer-aided analysis of second-order methods are [DKGT20], studying a single iteration of Newton's method on $\mathcal{S}_{M,+}$ and building on conditions that are only necessary for this class, and [Hil21], which obtained tight guarantees

on the same setting, but did not use interpolation conditions. More generally, all analysis, whether conducted by hand or in a computer-aided way, builds on necessary conditions and is hence a priori untight.

7.2.3 Univariate second-order optimization

Clearly, a computer-aided technique encompassing the analysis of multivariate second-order settings, as for first-order settings, would have been desirable. However, our proposed extension of the PEP framework to second-order analysis is restricted to univariate functions due to two main difficulties. First, the approach to second-order interpolation conditions presented in Chapter 6 heavily relies on properties specific to univariate functions. Second-order interpolation conditions for classes of multivariate functions remain an open question. Second, there is no known efficient way to solve PEPs involving second-order quantities in the multivariate case.

Nevertheless, tightly analyzing univariate second-order optimization is interesting in itself. Indeed, exact univariate bounds already provide insight into the bottlenecks of the method, and are lower bounds on the multivariate worst-case performance of the method. In fact, in many settings (first and second-order methods), univariate functions are observed to provide the general multivariate worst-case [CGT10, Doi21, THG17c, Toi24, RGP24]. In addition, in case the obtained univariate lower bound matches a known analytical multivariate upper bound, the associated worst-case instance proves tightness of this upper bound.

7.3 Interpolation of generalized self-concordant functions

Building on the approach proposed in Chapter 6, this section derives interpolation conditions (with or without function values) for \mathcal{H}_M , and the classes $\int^{(2)} \mathcal{F}_\alpha$ of univariate generalized self-concordant functions [STD19, Definition 1], satisfying

$$|f'''(x)| \leq M_\alpha f''(x)^\alpha, \quad f''(x) \geq 0, \quad (7.5)$$

where $\alpha \geq 0$. Letting $\alpha = 0$, $\alpha = 1$, and $\alpha = \frac{3}{2}$ allows recovering $\mathcal{H}_{M,0}$, $\mathcal{S}_{M,+}$, and $\mathcal{T}_{M,+}$, respectively. As this class $\int^{(2)} \mathcal{F}_\alpha$ encompasses functions whose

second derivative belongs to some class \mathcal{F}_α , it is amenable to the approach of Chapter 6, provided the assumptions required by Theorem 6.1 hold, and interpolation conditions for the associated *basic function classes* \mathcal{F}_α are known.

We thus first present \mathcal{F}_α and several characterizations thereof, before (i) obtaining interpolation conditions for \mathcal{F}_α , and (ii) iteratively applying Theorem 6.1, after verifying all assumptions required for its validity hold.

7.3.1 Interpolation of generalized Lipschitz functions

Basic classes \mathcal{F}_α are defined as satisfying a notion of generalized Lipschitz-ness, in the spirit of generalized smoothness introduced in [LQT⁺24] and generalized self-concordance introduced in [STD19]. Throughout, given $\alpha \geq 0$, we define

$$\beta(\alpha) = \begin{cases} \frac{1}{1-\alpha} & \text{if } \alpha \neq 1 \\ 1 & \text{if } \alpha = 1. \end{cases} \tag{7.6}$$

Definition 7.1 (Generalized Lipschitz function). Let $M, \alpha \geq 0$, and $\beta(\alpha)$ be defined as in (7.6). We say that $f : \mathbb{R} \rightarrow \bar{\mathbb{R}}$ is $(M, \alpha, +)$ -generalized Lipschitz if f (i) is non-negative everywhere, (ii) belongs to $\bar{\mathcal{C}}^0$ if $\alpha > 1$, and \mathcal{C}^0 if $\alpha \leq 1$, (iii) is piecewise $\bar{\mathcal{C}}^1$ if $\alpha > 1$, and piecewise \mathcal{C}^1 if $\alpha \leq 1$, and (iv) satisfies, whenever differentiable,

$$|f'(x)| \leq |\beta(\alpha)| M f(x)^\alpha. \tag{7.7}$$

We denote by $\mathcal{F}_{M, \alpha, +} \subseteq \bar{\mathcal{C}}^0$ (and \mathcal{C}^0 when $\alpha \leq 1$) the class of $(M, \alpha, +)$ -generalized Lipschitz functions. In addition, $f : \mathbb{R} \rightarrow \mathbb{R}$ is M -Lipschitz if (i) $f \in \mathcal{C}^0$, (ii) f is piecewise \mathcal{C}^1 if $\alpha \leq 1$, and (iii) f satisfies, whenever differentiable, (7.7).

We use the notation $\mathcal{F}_{M, \alpha, (+)}$ to refer to both $\mathcal{F}_{M, \alpha, +}$ and \mathcal{F}_M . Non negativity of $f(x)$ is required (i) as to describe *convex* function classes $\int^2 \mathcal{F}_{M, \alpha, +}$, that is, to fit (7.5), and (ii) as it simplifies the derivation of extremal interpolants for $\mathcal{F}_{M, \alpha, +}$, see Proposition 7.3.

Proposition 7.1 (2-points definition). Let $M, \alpha \geq 0$, $f \in \bar{\mathcal{C}}^0$, and f piecewise $\bar{\mathcal{C}}^1$. In addition, if $\alpha \leq 1$, let $f \in \mathcal{C}^0$, and f piecewise \mathcal{C}^1 . Then, $f \in \mathcal{F}_{M, \alpha, +}$, if and

only if $f = 0$, or $\forall x, y \in \mathbb{R}$:

$$|\tilde{f}(x) - \tilde{f}(y)| \leq M|x - y|, \tag{7.8}$$

$$\begin{cases} f(x) \geq 0, & \text{if } \alpha < 1, \\ f(x) > 0, & \text{if } \alpha \geq 1, \end{cases} \tag{7.9}$$

where $\tilde{f}(x) = \begin{cases} f(x)^{1-\alpha}, & \text{if } \alpha \neq 1, \\ \log(f(x)), & \text{if } \alpha = 1. \end{cases}$

In addition, $f \in \mathcal{F}_M$ if and only if it satisfies (7.8).

Proof. When $f'(x)$ exists and $f(x) \neq 0$, $\tilde{f}'(x) = \frac{f'(x)}{|\beta(\alpha)|f(x)^\alpha} \forall \alpha \geq 0$. In addition, on the interior of any interval over which $f = 0$, so are \tilde{f} and \tilde{f}' .

(Necessity of (7.8)) Let $f \in \mathcal{F}_{M,\alpha,+}$. By assumptions on f and definition of \tilde{f} , \tilde{f} is continuous and piecewise \mathcal{C}^1 . Indeed, if $\alpha \leq 1$, $f \in \mathcal{C}^0$ and is piecewise \mathcal{C}^1 , and \tilde{f} preserves these properties, with points of discontinuity the intersection of all points where f is non-differentiable, and the points where f switches from non-zero to zero values. On the other hand, if $\alpha > 1$, then $\tilde{f} = 0$ when $f = \infty$, hence, $\tilde{f} \in \mathcal{C}^0$ and is piecewise \mathcal{C}^1 despite $f \in \mathcal{C}^0$, and piecewise \mathcal{C}^1 .

Let $x_1 < x_2 < \dots < x_K$ be the points at which \tilde{f} is non-differentiable. Consider any interval \mathcal{I}_i on which \tilde{f} is differentiable, and non-zero. By the mean value theorem, it holds that $\forall x < y \in \mathcal{I}_i, \exists c \in (x, y)$ such that

$$\frac{|\tilde{f}(x) - \tilde{f}(y)|}{|x - y|} = |\tilde{f}'(c)| = \frac{|f'(c)|}{|\beta(\alpha)||f(c)^\alpha|} \stackrel{(7.7)}{\leq} \frac{|\beta(\alpha)|M|f(c)^\alpha|}{|\beta(\alpha)||f(c)^\alpha|} = M.$$

In addition, on any interval over which $\tilde{f} = 0$, \tilde{f} is M -Lipschitz. Therefore, \tilde{f} is piecewise M -Lipschitz. Consider now $x < x_i < \dots < x_j < y$ where $1 \leq i \leq j \leq K$; i.e., x, y may belong to distinct intervals. Then,

$$\begin{aligned} \frac{|\tilde{f}(x) - \tilde{f}(y)|}{|x - y|} &\leq \frac{|\tilde{f}(x) - \tilde{f}(x_i)| + |\tilde{f}(x_i) - \tilde{f}(x_{i+1})| + \dots + |\tilde{f}(x_j) - \tilde{f}(y)|}{|x - y|} \\ &\leq M \frac{x_i - x + x_{i+1} - x_i + \dots + y - x_j}{y - x} = M. \end{aligned}$$

Finally, nonnegativity is due to Definition 7.1 and positivity in the case $\alpha \geq 1$ arises from the limit case of (7.8): if $f \in \mathcal{F}_{M,\alpha,+}$ and $f(x) = 0$ at some $x \in \mathbb{R}$, then $f = 0$.

(Sufficiency of (7.8)) Let $f \in \mathcal{C}^0$, piecewise \mathcal{C}^1 and satisfying (7.8). At all $x \in \mathbb{R}$ where f is differentiable and non-zero, it holds:

$$|\tilde{f}'(x)| = \frac{|f'(x)|}{|\beta(\alpha)||f(x)|^\alpha} = \lim_{h \rightarrow 0} \frac{|\tilde{f}(x+h) - \tilde{f}(x)|}{h} \stackrel{(7.8)}{\leq} M$$

$$\Rightarrow |f'(x)| \leq |\beta(\alpha)|Mf(x)^\alpha.$$

In addition, at all x where $f(x) = 0$, by nonnegativity of f , either $f'(x)$ does not exist, or $f'(x) = 0$. Hence, $f \in \mathcal{F}_{M,\alpha,+}$.

Finally, the case \mathcal{F}_M follows the same argument, except there is no need to handle the non-negativity constraint. \square

Proposition 7.1 reduces generalized self-concordant functions to functions f whose associated quantity \tilde{f}'' is simply Lipschitz continuous, which allows obtaining interpolation conditions almost straightforwardly for all basic function classes and their higher-order associated classes, once interpolation conditions are known for \mathcal{F}_M and $\int^{(k)} \mathcal{F}_M, k = 1, 2, \dots$. In addition, the conditions in Proposition 7.1 are interpolation conditions for $\mathcal{F}_{M,\alpha,(+)}$.

Theorem 7.1 (Interpolation conditions for $\mathcal{F}_{M,\alpha,(+)}$)

Let $M, \alpha \geq 0$. A set $S = \{(x_i, f_i)\}_{i \in [N]} \in (\mathbb{R} \times \mathbb{R})^N$ is $\mathcal{F}_{M,\alpha,+}$ -interpolable if and only if $f_i = 0, \forall i \in [N]$, or $\forall i, j \in [N]$:

$$|\tilde{f}_i - \tilde{f}_j| \leq M|x_i - x_j|, \tag{7.10}$$

$$\begin{cases} f_i \geq 0 & \text{if } \alpha < 1, \\ f_i > 0 & \text{if } \alpha \geq 1, \end{cases} \tag{7.11}$$

where $\tilde{f}_i = \begin{cases} f_i^{1-\alpha}, & \text{if } \alpha \neq 1, \\ \log(f_i), & \text{if } \alpha = 1. \end{cases}$

In addition, S is \mathcal{F}_M -interpolable if and only if it satisfies (7.10).

Proof. By Proposition 7.1, these conditions are necessary for interpolation. We show they are also sufficient. Suppose S satisfies (7.10). We construct a function $f \in \mathcal{F}_{M,\alpha,(+)}$ interpolating S . Let $\tilde{f}(x) = \min_k \tilde{f}_k + M|x - x_k|$ and

$$f(x) = \begin{cases} \tilde{f}(x)^{\alpha-1} & \text{if } \alpha \neq 1, \\ e^{\tilde{f}(x)} & \text{if } \alpha = 1. \end{cases} \tag{7.12}$$

Since, by (7.10), $\tilde{f}_i \leq \tilde{f}_j + M|x_j - x_i|$, $\forall i, j \in [N]$, $f(x)$ as defined in (7.12) interpolates S , i.e.,

$$\text{when } \alpha \neq 1 : f(x_i) = (\min_k \tilde{f}_k + M|x_i - x_k|)^{\alpha-1} = \tilde{f}_i^{\alpha-1} = f_i.$$

$$\text{when } \alpha = 1 : f(x_i) = e^{\min_k \tilde{f}_k + M|x_i - x_k|} = e^{\tilde{f}_i} = f_i$$

Moreover, $f \in \mathcal{F}_{M,\alpha,+}$ since $\forall x, y \in \mathbb{R}$, we have

$$\begin{aligned} |\tilde{f}(x) - \tilde{f}(y)| &= |\tilde{f}_{k_*(x)} + M|x - x_{k_*(x)}| - \tilde{f}_{j_*(y)} - M|y - x_{j_*(y)}|| \\ &\leq |\tilde{f}_{j_*(y)} + M|x - x_{j_*(y)}| - \tilde{f}_{j_*(y)} - M|y - x_{j_*(y)}|| \\ &= M||x - x_{j_*(y)}| - |y - x_{j_*(y)}|| \leq M|x - y|, \end{aligned}$$

where $k_*(x)$ and $j_*(y)$ are the optimal indices in the definition of $\tilde{f}(x)$ and $\tilde{f}(y)$. The last inequality follows from the reverse triangle inequality. In addition, if $f_i \geq 0 \forall i \in [N]$ then $f(x) \geq 0 \forall x \in \mathbb{R}$, and the same holds if $f_i > 0$. Finally, $\tilde{f}(x) \in \mathcal{C}^0$ and piecewise \mathcal{C}^1 . \square

Remark 7.1. By Lemma 6.1, Theorem 7.1 furnishes interpolation conditions without function values for $\int \mathcal{F}_{M,\alpha,(+)}$, and interpolation conditions without function values and first derivative for $\int^{(2)} \mathcal{F}_{M,\alpha,(+)}$.

Corollaries of Theorem 7.1 It turns out that several particular cases of Theorem 7.1 recover well-known conditions. For instance,

- Interpolation conditions for \mathcal{F}_M are the classical interpolation conditions for Lipschitz continuous functions, see Section 6.2.
- Setting $\alpha = \frac{3}{2}$ yields interpolation conditions without function and gradient values, for the class of self-concordant functions, i.e.

$$\left| \frac{1}{\sqrt{h_i}} - \frac{1}{\sqrt{h_j}} \right| \leq M|x_j - x_i|.$$

This condition, the strongest possible that only involves second-order derivatives and points, was derived in [NN94, Theorem 2.1.1] in the multivariate case.

- Letting $\alpha = 1$ yields interpolation conditions without function and gradient values, for the class of quasi-self-concordant functions. This

condition, given by

$$|\log(g_i) - \log(g_j)| \leq M|x_i - x_j|,$$

has been derived in [Doi23, Lemma 2.5], but also in [VRN⁺24, Lemma 2.5] regarding the class of (L_0, L_1) -smooth functions [ZHS]19], with $L_0 = 0$ and $L_1 = M$.

7.3.2 Interpolation of generalized smooth functions

To derive interpolation conditions for $\int \mathcal{F}_{M,\alpha,(+)}$ via the approach proposed in Chapter 6, we show that $\mathcal{F}_{M,\alpha,(+)}$ and $\int \mathcal{F}_{M,\alpha,(+)}, k = 1, 2$, satisfy the assumptions in Theorem 6.1. We first consider assumptions that are independent of extremal interpolants of $\mathcal{F}_{M,\alpha,(+)}$, i.e., connectability and convexity.

Proposition 7.2 (Properties of generalized Lipschitz functions). *Let $\alpha, M \geq 0$. Then, $\mathcal{F}_{M,\alpha,(+)}$, $\int \mathcal{F}_{M,\alpha,(+)}$ and $\int^{(2)} \mathcal{F}_{M,\alpha,(+)}$ are respectively order 0, 1 or 2 connectable (Assumption 6.4). In addition, when $\alpha \leq 1$, $\int \mathcal{F}_{M,\alpha,(+)}$ and $\int^{(2)} \mathcal{F}_{M,\alpha,(+)}$ are also convex (Assumption 6.3).*

Proof. (Connectability) Definition 7.1, involving derivatives, is pointwise and holds everywhere except for a finite set of points. Hence, by definition, $\mathcal{F}_{M,\alpha,(+)}$ is order 0-connectable, and the function f , juxtaposition of several functions in $\mathcal{F}_{M,\alpha,(+)}$ might be non-differentiable only at the junctions between intervals. The same argument holds for $\int \mathcal{F}_{M,\alpha,(+)}$ and $\int^{(2)} \mathcal{F}_{M,\alpha,(+)}$, except that it involves higher-order derivatives.

(Convexity) Let $\alpha \leq 1, f_1, f_2 \in \int \mathcal{F}_{M,\alpha,(+)}, \lambda \in [0, 1]$, and $f = \lambda f_1 + (1 - \lambda) f_2$. It holds:

$$\begin{aligned} |f''(x)| &= |\lambda f_1''(x) + (1 - \lambda) f_2''(x)| \leq \lambda |f_1''(x)| + (1 - \lambda) |f_2''(x)| \\ &\stackrel{(7.7)}{\leq} M_1 (\lambda f_1'(x)^\alpha + (1 - \lambda) f_2'(x)^\alpha) \leq M_1 (\lambda f_1'(x) + (1 - \lambda) f_2'(x))^\alpha \end{aligned} \tag{7.13}$$

by concavity of $g(t) = t^\alpha, t \geq 0, \alpha \leq 1$. Moreover, $\forall x \in \mathbb{R}$, if $f_1(x), f_2(x) \geq 0$ then $f(x) \geq 0$. Hence, $f \in \int \mathcal{F}_{M,\alpha,0}$, which is convex. The same argument holds for $\int^{(2)} \mathcal{F}_{M,\alpha,(+)}$ (with f''' and f'' instead of f'' and f' in (7.13)). \square

We now show extremal interpolability of $\mathcal{F}_{M,\alpha,(+)}$, and provide its extremal interpolants.

Proposition 7.3 (Extremal interpolants of $\mathcal{F}_{M,\alpha,(+)}$). *Let $\alpha, M \geq 0$. Then, $\mathcal{F}_{M,\alpha,(+)}$ is extremally interpolable (Assumption 6.1), and the extremal interpolants (Definition 6.3) of an $\mathcal{F}_{M,\alpha,(+)}$ -interpolable set $S = \{(x_i, f_i)\}_{i=1,2} \in (\mathbb{R} \times \mathbb{R})^2$, where $x_1 < x_2$, are given by:*

- [f_{\min} , Case 1] *If $\alpha \geq 1$, if we consider \mathcal{F}_M , or if $\alpha < 1$ and $\tilde{f}_2 + \tilde{f}_1 \geq M(x_2 - x_1)$:*

$$f_{\min}(x) = \begin{cases} \nu \left(\tilde{f}_1 - \frac{\beta(\alpha)}{|\beta(\alpha)|} M(x - x_1) \right) & x \in [x_1, z] \\ \nu \left(\tilde{f}_2 + \frac{\beta(\alpha)}{|\beta(\alpha)|} M(x - x_2) \right) & x \in [z, x_2] \end{cases}'$$

where $z = \frac{x_1+x_2}{2} + \frac{\beta(\alpha)}{|\beta(\alpha)|} \frac{\tilde{f}_1 - \tilde{f}_2}{2M}$.

- [f_{\min} , Case 2] *Else:*

$$f_{\min}(x) = \begin{cases} \nu \left(\tilde{f}_1 - \frac{\beta(\alpha)}{|\beta(\alpha)|} M(x - x_1) \right) & x \in [x_1, x_1 + \frac{\tilde{f}_1}{M}] \\ 0 & x \in [x_1 + \frac{\tilde{f}_1}{M}, x_2 - \frac{\tilde{f}_2}{M}] \\ \nu \left(\tilde{f}_2 + \frac{\beta(\alpha)}{|\beta(\alpha)|} M(x - x_2) \right) & x \in [x_2 - \frac{\tilde{f}_2}{M}, x_2] \end{cases}.$$

- [f_{\max} , Case 1] *If $\alpha \leq 1$, or if $\alpha > 1$ and $\tilde{f}_2 + \tilde{f}_1 > M(x_2 - x_1)$:*

$$f_{\max}(x) = \begin{cases} \nu \left(\tilde{f}_1 + \frac{\beta(\alpha)}{|\beta(\alpha)|} M(x - x_1) \right) & x \in [x_1, y] \\ \nu \left(\tilde{f}_2 - \frac{\beta(\alpha)}{|\beta(\alpha)|} M(x - x_2) \right) & x \in [y, x_2] \end{cases}'$$

where $y = \frac{x_1+x_2}{2} - \frac{\beta(\alpha)}{|\beta(\alpha)|} \frac{\tilde{f}_1 - \tilde{f}_2}{2M}$.

- [f_{\max} , Case 2] *Else:*

$$f_{\max}(x) = \begin{cases} \nu \left(\tilde{f}_1 + \frac{\beta(\alpha)}{|\beta(\alpha)|} M(x - x_1) \right) & x \in [x_1, x_1 + \frac{\tilde{f}_1}{M}] \\ \infty & x \in [x_1 + \frac{\tilde{f}_1}{M}, x_2 - \frac{\tilde{f}_2}{M}] \\ \nu \left(\tilde{f}_2 - \frac{\beta(\alpha)}{|\beta(\alpha)|} M(x - x_2) \right) & x \in [x_2 - \frac{\tilde{f}_2}{M}, x_2] \end{cases},$$

where

$$\begin{aligned} \tilde{f}_i &= f_i^{1-\alpha} && \text{if } \alpha \neq 1, && \log(f_i) \text{ otherwise,} \\ \nu(x) &= x^{\beta(\alpha)} && \text{if } \alpha \neq 1, && e^x \text{ otherwise.} \end{aligned}$$

Proof. First, observe that all quantities in Proposition 7.3 are well-defined, i.e., $x_1 \leq z \leq x_2$ and $x_1 \leq y \leq x_2$. Indeed, e.g.,

$$z - x_1 = \frac{x_2 - x_1}{2} + \frac{\beta(\alpha)}{|\beta(\alpha)|} \frac{\tilde{f}_1 - \tilde{f}_2}{2M} \geq 0$$

by satisfaction of (7.8).

We now show that f_{\min}, f_{\max} , as illustrated on Figures 7.1 and 7.2, are the extremal interpolating envelopes of S as defined in (6.7). It holds that f_{\min}, f_{\max} interpolate S , since $v(\tilde{f}_i) = f_i, i = 1, 2$. For any $x \in [x_1, x_2]$, they are also the extremal f satisfying (7.10) with respect to S , i.e.,

$$f_{\min(\text{ resp. } \max)}(x) = \min(\text{ resp. } \max)_f f \text{ s.t. } |\tilde{f} - \tilde{f}_i| \leq M|x - x_i|, i = 1, 2, \\ f \geq 0 \text{ (except for } \mathcal{F}_M).$$

Case $\alpha < 1$ (Figure 7.1a). Since $\beta(\alpha) > 0$, $v(x)$ is an increasing nonnegative function of x , when $x \geq 0$, and either (i) $v(x)$ does not exist for $x < 0$ (e.g., $\alpha = \frac{1}{3}$), (ii) $v(x)$ is negative for $x < 0$ (e.g., $\alpha = 0$), or (iii) $v(x)$ is a decreasing nonnegative function for $x < 0$ (e.g., $\alpha = \frac{1}{2}$).


Hence, at all $x \in [x_1, x_2]$, computing the extremal interpolants of S , i.e., $f_{\min(\text{ resp. } \max)}$, is equivalent to computing the extremal M -Lipschitz interpolants of $\{(x_i, \tilde{f}_i)\}_{i=1,2}$, whenever these are nonnegative. On the contrary, when these interpolants become negative, the extremal interpolants of S are set to 0.

Specifically, at all $x \in [x_1, x_2]$, $\tilde{f}_{\min}(x)$ takes the maximal value between 0 and the lower boundaries $\tilde{f} \geq \tilde{f}_i - M|x - x_i|$. If $\tilde{f}_1 + \tilde{f}_2 \geq M(x_2 - x_1)$, these boundaries are always larger than 0, and f_{\min} consists in two pieces. Else, f_{\min} consists in three parts, since the associated function $\tilde{f}_{\min}(x)$ is set to 0 whenever the lower boundaries $\tilde{f}_i - M|x - x_i|$ become negative. In addition, at all $x \in [x_1, x_2]$, $\tilde{f}_{\max}(x)$ takes the minimal value of the upper boundaries $\tilde{f} \leq \tilde{f}_i + M|x - x_i|$, and always consists in two parts, since these boundaries are always nonnegative.


Case $\alpha = 1$ (Figure 7.1b). Since $v(x)$ is an everywhere increasing nonnegative function, computing the extremal interpolants of S , i.e., $f_{\min(\text{ resp. } \max)}$, is equivalent to computing the extremal M -Lipschitz interpolants of $\{(x_i, \tilde{f}_i)\}_{i=1,2}$, i.e., $\tilde{f}_{\min(\text{ resp. } \max)}$.

Specifically, at all $x \in [x_1, x_2]$, $\tilde{f}_{\min}(x)$ takes the maximal value between the lower boundaries $\tilde{f} \geq \tilde{f}_i - M|x - x_i|$, even when such boundaries become negative. Similarly, at all $x \in [x_1, x_2]$, $\tilde{f}_{\max}(x)$ takes the minimal value

of the upper boundaries $\tilde{f} \leq \tilde{f}_i + M|x - x_i|$. Both extremal interpolants always consist of two parts. Case $\alpha > 1$ (Figures 7.2a and 7.2b). Since $\beta(\alpha) < 0$,



(a) Illustration of the extremal interpolants for $\alpha < 1$, on $\alpha = 0$. The left-hand side figure considers \tilde{S}_1 , and the associated extremal interpolants consist of 2 pieces, i.e., $\tilde{f}_{\min}(x)$ is naturally always nonnegative. The right-hand side figure considers \tilde{S}_2 , and the associated minimal interpolant consists in 3 pieces, i.e., $\tilde{f}_{\min}(x)$ is forced to be zero when the M -Lipschitz minimal interpolant of \tilde{S}_2 is negative.




(b) Extremal interpolants for $\alpha = 1$, on \tilde{S}_2 . The right-hand side figure displays \tilde{f}_{\min} and \tilde{f}_{\max} , the extremal M -Lipschitz interpolants of \tilde{S}_2 , even when negative. The left-hand side figure displays f_{\min} and f_{\max} , the actual extremal interpolants of S_2 , that are always nonnegative.

Fig. 7.1 Given $\alpha \geq 0$, $M = 1$, and $S = \{(x_i, f_i)\}_{i=1,2}$, extremal interpolants $f_{\min}(x)$, $f_{\max}(x)$ (Definition 6.3) of S , depending on α . The considered sets are either $\tilde{S}_1 = \{(x_i, \tilde{f}_i)\}_{i=1,2} = \{(0, \frac{1}{2}), (1, \frac{3}{4})\}$, or $\tilde{S}_2 = \{(x_i, \tilde{f}_i)\}_{i=1,2} = \{(0, \frac{1}{2}), (1, \frac{1}{4})\}$.


$v(x)$ is a decreasing positive function of x , when $x > 0$, with asymptote in $x = 0$, and either (i) $v(x)$ does not exist for $x < 0$ (e.g., $\alpha = 3$), (ii) $v(x)$ is negative for $x < 0$ (e.g., $\alpha = 2$), or (iii) $v(x)$ is an increasing positive function for $x < 0$ (e.g., $\alpha = \frac{3}{2}$).

Specifically, at all $x \in [x_1, x_2]$, $\tilde{f}_{\min}(x)$ takes the minimal value of the upper boundaries $\tilde{f} \leq \tilde{f}_i + M|x - x_i|$ and always consists in two parts, since these boundaries are always positive. In addition, at all $x \in [x_1, x_2]$,

$\tilde{f}_{\max}(x)$ takes the maximal value between 0 and the lower boundaries $\tilde{f} \geq \tilde{f}_i - M|x - x_i|$. If $\tilde{f}_1 + \tilde{f}_2 > M(x_2 - x_1)$, these boundaries are always strictly larger than 0, and f_{\max} consists in two pieces, i.e., the associated functions to these boundaries cross each other before meeting their respective asymptote. Else, f_{\max} consists in three parts, since the associated function $\tilde{f}_{\max}(x)$ is set to 0 whenever the lower boundaries $\tilde{f}_i - M|x - x_i|$ become negative. On this interval, f_{\max} reaches $+\infty$.



(a) Extremal interpolants for $\alpha > 1$, on $\alpha = \frac{3}{2}$ and \tilde{S}_1 . The right-hand side figure displays \tilde{f}_{\min} and \tilde{f}_{\max} , the inverted extremal M -Lipschitz interpolants of \tilde{S}_1 , that are naturally nonnegative. The left-hand side figure displays f_{\min} and f_{\max} , the actual extremal interpolants of S_1 .



(b) Extremal interpolants for $\alpha > 1$, on $\alpha = \frac{3}{2}$ and \tilde{S}_2 . The right-hand side figure displays \tilde{f}_{\min} and \tilde{f}_{\max} , the inverted extremal M -Lipschitz interpolants of \tilde{S}_2 , forced to be zero when these become nonnegative. The left-hand side figure displays f_{\min} and f_{\max} , the actual extremal interpolants of S_2 .

Fig. 7.2 Follow up of Figure 7.1

To conclude the proof, it remains to show that $\mathcal{F}_{M,\alpha,(+)}$ is extremally interpolable, that is, $f_{\min}, f_{\max} \in \mathcal{F}_{M,\alpha,(+)}$. This holds since (i) f_{\min} and f_{\max} are nonnegative (and positive for $\alpha > 1$) except for the extremal interpolants of \mathcal{F}_M , (ii) f_{\min} and f_{\max} belong to $\bar{\mathcal{C}}^0$ (or \mathcal{C}^0 if $\alpha \leq 1$), and are piecewise $\bar{\mathcal{C}}^1$ (or \mathcal{C}^1 when $\alpha \leq 1$), and (iii) $\mathcal{F}_{M,\alpha,(+)}$ is order 0-connectable

and f_{\min}, f_{\max} are functions by parts satisfying on each interval:

$$|v(\tilde{f}_i \pm M(x - x_i))'| = M|\beta(\alpha)|v(\tilde{f}_i \pm M(x - x_i))^\alpha \quad i = 1, 2.$$

□

Remark 7.2. The cases \mathcal{F}_M and $\mathcal{F}_{M,1,\alpha}$ are the only ones for which the extremal interpolants take a single expression.

We build on these extremal interpolants to prove $\mathcal{F}_{M,\alpha,(+)}$ is extremally completable when $\alpha > 1$ (the case $\alpha \leq 1$ follows from Proposition 7.2).

Proposition 7.4 (Extremal completable of $\int \mathcal{F}_{M,\alpha,(+)}$). *Let $\alpha, M \geq 0$, and $\alpha > 1$. Then, $\int \mathcal{F}_{M,\alpha,(+)}$ is extremally completable.*

Proof. The proof is illustrated in Figure 7.3. Let $S = \{(x_i, f_i^0, f_i^1)\}_{i=1,2}$ and




Fig. 7.3 Illustration of the proof of Proposition 7.4. Given $S = \{(x_i, f_i, g_i)\}_{i=1,2}$, $\tilde{S} = \{(x_i, g_i)\}_{i=1,2}$, and its extremal interpolants $g_{\min}(x)$ and $g_{\max}(x)$, we construct a function $g \in \mathcal{F}_{M,\alpha,+}$, interpolating \tilde{S} , and whose intergral $f_1 + \int_{x_1}^{x_2} g(z, a) dz$ interpolates between $f_1 + \int_{x_1}^{x_2} g_{\min}(z) dz$ and $f_1 + \int_{x_1}^{x_2} g_{\max}(z) dz$. Specifically, we continuously decrease the integral of g_{\max} by 'cutting' g_{\max} at height a .

$g_{\min/\max}$ be defined as in (6.7), interpolating $\tilde{S} = \{(x_i, f_i^1)\}_{i=1,2}$. Let $a \geq 0$ and define the following function

$$g(x, a) = \min\{g_{\max}(x), \max\{g_{\min}(x), a\}\}.$$

By construction, g is defined as the juxtaposition of functions in $\mathcal{F}_{M,\alpha,+}$ on at most 3 intervals: either a constant function, either $g_{\min}(x)$ or $g_{\max}(x)$.

Hence, since $\mathcal{F}_{M,\alpha,+}$ is order 0-connectable, $g \in \mathcal{F}_{M,\alpha,+}$. In addition, g interpolates \tilde{S} since $g_{\max}(x_i) = g_{\min}(x_i) = f_i^1, i = 1, 2$. Let

$$f(x, a) := f_1^0 - \int_{-\infty}^{x_1} g(z, a) dz + \int_{-\infty}^x g(z, a) dz.$$

By Lemma 6.2, $f \in \int \mathcal{F}_{M,\alpha,+}$, $f'(x_i, a) = f_i^1, i = 1, 2$, and $f(x_1, a) = f_1^0$. In addition,

$$f(x_2, a) = f_1^0 + \int_{x_1}^{x_2} g(z, a) dz.$$

Since $f_1^0 + \int_{x_1}^{x_2} g_{\min}(z) dz \leq f_2^0 \leq f_1^0 + \int_{x_1}^{x_2} g_{\max}(z) dz$ and $\int_{x_1}^{x_2} g(z, a) dz$ depends continuously on a , with a minimum in $\int_{x_1}^{x_2} g_{\min}(z) dz$ and a maximum in $\int_{x_1}^{x_2} g_{\max}(z) dz$, there exists some $a \geq 0$ for which $f(x_2, a) = f_2^0$, hence $\int \mathcal{F}_{M,\alpha,+}$ is extremally completable. \square

Combining Propositions 7.4, 7.3, and 7.2, and Theorem 6.1, we obtain interpolation conditions for $\int \mathcal{F}_{M,\alpha,(+)}$. For the sake of clarity, we consider the case $\alpha = 1$ separately.

Theorem 7.2 (Interpolation conditions for $\int \mathcal{F}_{M,1,+}$)

Let $M \geq 0$. A set $S = \{(x_i, f_i, g_i)\}_{i \in [N]} \in (\mathbb{R} \times \mathbb{R} \times \mathbb{R})^N$ is $\int \mathcal{F}_{M,1,(+)}$ -interpolable if and only if, $\forall i, j \in [N], g_i \geq 0$ and

$$f_j - f_i \geq \frac{1}{M}(g_i + g_j) - \frac{2}{M} \sqrt{g_i g_j} e^{-\frac{M}{2}(x_j - x_i)}. \tag{7.14}$$

Proof of Theorem 7.2. Classes $\mathcal{F}_{M,1,(+)}$ and $\int \mathcal{F}_{M,1,(+)}$ satisfy all assumptions in Theorem 6.1. Hence, S is interpolable if and only if it satisfies

$$|\log(g_i) - \log(g_j)| \leq M|x_i - x_j|, \tag{7.15}$$

and $\forall x_i < x_j$:

$$\begin{aligned} & \begin{cases} f_j - f_i \geq \frac{g_i + g_j}{M} - \frac{2}{M} \sqrt{g_i g_j} e^{-\frac{M}{2}(x_j - x_i)} \\ f_j - f_i \leq -\frac{g_i + g_j}{M} + \frac{2}{M} \sqrt{g_i g_j} e^{\frac{M}{2}(x_j - x_i)} \end{cases} \\ \Leftrightarrow & f_i - f_j \leq \frac{g_i + g_j}{M} - \frac{2}{M} \sqrt{g_i g_j} e^{-\frac{M}{2}(x_i - x_j)}. \end{aligned} \tag{7.16}$$

One can show that imposing (7.14) on all pairs is strictly equivalent to (7.15) and (7.16). In addition, (7.14) implies (7.15) when applied to all pairs. Indeed, summing (7.14) imposed on (i, j) and (j, i) yields

$$\frac{g_i + g_j}{\sqrt{g_i g_j}} \leq e^{-\frac{M}{2}(x_j - x_i)} + e^{\frac{M}{2}(x_j - x_i)}$$

and using $\frac{1}{t} + t \leq \frac{1}{\sqrt{s}} + \sqrt{s} \Leftrightarrow \frac{1}{s} \leq t^2 \leq s$ when $s \geq 1, t \geq 0$ with $t = \sqrt{\frac{g_j}{g_i}}$ and $s = e^{M(x_j - x_i)}$ yields

$$e^{-M(x_j - x_i)} \leq \frac{g_j}{g_i} \leq e^{M(x_j - x_i)} \tag{7.17}$$

which is equivalent to (7.15). □

Theorem 7.3 (Interpolation conditions for $\int \mathcal{F}_{M,\alpha,+}$ ($\alpha \neq 1$))

Let $\alpha \neq 1, M \geq 0$. A set $S = \{(x_i, f_i, g_i)\}_{i \in [N]} \in (\mathbb{R} \times \mathbb{R} \times \mathbb{R})^N$ is $\int \mathcal{F}_{M,\alpha,+}$ -interpolable if and only if, $\forall i, j \in [N], g_i = 0$ and $f_i = f_j$, or $\forall i, j \in [N]$:

$$|\tilde{g}_i - \tilde{g}_j| < M|x_i - x_j|, \text{ and } g_i \geq 0 \text{ if } \alpha < 1, g_i > 0 \text{ else,} \tag{7.18}$$

If $\tilde{g}_i + \tilde{g}_j \geq \frac{\beta(\alpha)}{|\beta(\alpha)|} M(x_j - x_i)$:

$$f_j - f_i \geq \frac{\beta(\alpha)}{|\beta(\alpha)| M(\beta(\alpha) + 1)} \left(\tilde{g}_i^{\beta(\alpha)+1} + \tilde{g}_j^{\beta(\alpha)+1} - \frac{1}{2^{\beta(\alpha)}} (\tilde{g}_i + \tilde{g}_j - \frac{\beta(\alpha)}{|\beta(\alpha)|} M(x_j - x_i))^{\beta(\alpha)+1} \right), \tag{7.19}$$

If $\alpha < 1$ and $\tilde{g}_i + \tilde{g}_j \leq M(x_j - x_i)$:

$$f_j - f_i \geq \frac{1}{M(\beta(\alpha) + 1)} \left(\tilde{g}_i^{\beta(\alpha)+1} + \tilde{g}_j^{\beta(\alpha)+1} \right). \tag{7.20}$$

where $\tilde{g}_i = g_i^{1/\beta(\alpha)}$. In addition, S is $\int \mathcal{F}_M$ -interpolable if and only if it satisfies (7.19).

Proof of Theorem 7.3. Classes $\mathcal{F}_{M,\alpha,+}$ and $\int \mathcal{F}_{M,\alpha,+}$ satisfy all assumptions in Theorem 6.1. Hence, S is interpolable if and only if it satisfies (7.18), and

$\forall x_i < x_j$:

- Integration of $f_{\min}(x)$, Case 1:

If $\alpha > 1$, or if $\alpha < 1$ and $\tilde{g}_i + \tilde{g}_j \geq M(x_j - x_i)$, or considering $\mathcal{F}_{M,0}$:

$$f_j - f_i \geq \frac{\beta(\alpha)}{|\beta(\alpha)|M(\beta(\alpha) + 1)} \left(\tilde{g}_i^{\beta(\alpha)+1} + \tilde{g}_j^{\beta(\alpha)+1} - \frac{1}{2^{\beta(\alpha)}} \left(\tilde{g}_i + \tilde{g}_j - \frac{\beta(\alpha)}{|\beta(\alpha)|} M(x_j - x_i) \right)^{\beta(\alpha)+1} \right)$$

- Integration of $f_{\min}(x)$, Case 2: If $\alpha < 1$ and $\tilde{g}_i + \tilde{g}_j \leq M(x_j - x_i)$:

$$f_j - f_i \geq \frac{1}{M(\beta(\alpha) + 1)} \left(\tilde{g}_i^{\beta(\alpha)+1} + \tilde{g}_j^{\beta(\alpha)+1} \right)$$

- Integration of $f_{\max}(x)$, Case 1: If $\alpha < 1$, or if $\alpha > 1$ and

$$\tilde{g}_i + \tilde{g}_j > M(x_j - x_i) :$$

$$f_j - f_i \leq \frac{\beta(\alpha)}{|\beta(\alpha)|M(\beta(\alpha) + 1)} \left(-\tilde{g}_i^{\beta(\alpha)+1} - \tilde{g}_j^{\beta(\alpha)+1} + \frac{1}{2^{\beta(\alpha)}} \left(\tilde{g}_i + \tilde{g}_j + \frac{\beta(\alpha)}{|\beta(\alpha)|} M(x_j - x_i) \right)^{\beta(\alpha)+1} \right)$$

$$\Leftrightarrow f_i - f_j \geq \frac{\beta(\alpha)}{|\beta(\alpha)|M(\beta(\alpha) + 1)} \left(\tilde{g}_i^{\beta(\alpha)+1} + \tilde{g}_j^{\beta(\alpha)+1} - \frac{1}{2^{\beta(\alpha)}} \left(\tilde{g}_i + \tilde{g}_j - \frac{\beta(\alpha)}{|\beta(\alpha)|} M(x_i - x_j) \right)^{\beta(\alpha)+1} \right).$$

The second case for f_{\max} , i.e., when it reaches infinite values, yields no further constraint. Considering $\mathcal{F}_{M,\alpha,+}$, imposing (7.19) and (7.20) on all pairs is strictly equivalent to these conditions, since $\frac{\beta(\alpha)}{|\beta(\alpha)|} = -1$ if $\alpha > 1$, $\frac{\beta(\alpha)}{|\beta(\alpha)|} = 1$ otherwise. Considering $\mathcal{F}_{M,0}$, there are no subcases to consider, and imposing (7.19) on all pairs is strictly equivalent to these conditions. Furthermore, satisfaction of (7.19) on both pairs (i, j) and (j, i) implies satisfaction of $|g_i - g_j| \leq M|x_j - x_i|$, since adding both inequalities gives

$$0 \geq g_i^2 + g_j^2 - \frac{1}{4}(g_i + g_j - M(x_j - x_i))^2 - \frac{1}{4}(g_i + g_j + M(x_j - x_i))^2$$

$$\Leftrightarrow 0 \geq \frac{(g_i - g_j)^2}{2} - \frac{M^2}{2}(x_j - x_i)^2 \Leftrightarrow |g_i - g_j| \leq M|x_i - x_j|.$$

□

Remark 7.3. By Proposition 6.1, Theorems 7.2 and 7.3 furnish interpolation conditions without function values for $\int^{(2)} \mathcal{F}_{M,\alpha,(+)}$.

Remark 7.4. Necessarily, the interpolation conditions of Theorems 7.2 and 7.3 are satisfied everywhere by any function in $\int \mathcal{F}_{M,\alpha,(+)}$. They are thus equivalent to Definition 4.1 of $\int \mathcal{F}_{M,\alpha,(+)}$ when imposed everywhere, even though proving it from scratch is not straightforward.

7.3.3 Applications of Theorems 7.2 and 7.3

We propose a list of corollaries to Theorems 7.2 and 7.3, specific several specific classes $\int^{(2)} \mathcal{F}_{M,\alpha,(+)}$.

Class \mathcal{H}_M of functions with M -Lipschitz Hessian We recover the interpolation conditions for M -smooth functions from [THG17c, Theorem 4].

Corollary 7.1. *A set $\{(x_i, f_i, g_i, h_i)\}_{i \in [N]}$ is \mathcal{H}_M -interpolable without functions values if, and only if, $\forall i, j \in [N]$:*

$$g_j - g_i - h_i(x_j - x_i) \geq -\frac{M}{2}(x_j - x_i)^2 + \frac{1}{4M}(h_j - h_i + M(x_j - x_i))^2. \quad (7.21)$$

Proof. Since $\mathcal{H}_M = \int^{(2)} \mathcal{F}_M$, consider Theorem 7.3 as applied to $\int \mathcal{F}_M$. □

Class $\mathcal{H}_{M,\mu}$ of μ -strongly convex functions with M -Lipschitz Hessian

Corollary 7.2. *A set $\{(x_i, f_i, g_i, h_i)\}_{i \in [N]} \in (\mathbb{R} \times \mathbb{R} \times \mathbb{R} \times \mathbb{R})^N$ is $\mathcal{H}_{M,\mu}$ -interpolable without functions values if and only, $\forall i, j \in [N]$:*

$$g_j - g_i \geq h_i(x_j - x_i) - \frac{M}{2}(x_j - x_i)^2 + \frac{1}{4M}(h_j - h_i + M(x_j - x_i))^2 \quad (7.22)$$

$$h_i \geq \mu \quad (7.23)$$

$$\text{If } x_j - x_i \geq \frac{h_i + h_j - 2\mu}{M}, \text{ then } g_j - g_i \geq \mu(x_j - x_i) + \frac{(h_i - \mu)^2 + (h_j - \mu)^2}{2M}. \quad (7.24)$$

Proof. When $\mu = 0$, $\mathcal{H}_{M,\mu} = \int^{(2)} \mathcal{F}_{M,0,+}$, hence Corollary 7.2 is an application of Theorem 7.3 as applied to $\mathcal{F}_{M,0,+}$. The case $\mu \geq 0$ follows from the observation that $f \in \int^{(2)} \mathcal{F}_{M,0,+} \Leftrightarrow f + \frac{\mu}{2}x^2 \in \mathcal{H}_{M,\mu}$. \square

Remark 7.5. Condition (7.22) is the interpolation condition for M -smoothness of S , while (7.23) (lower bound on h) ensures S to be consistent with an increasing function, of slope at least μ . However, juxtaposing these conditions alone is not an interpolation condition for the class of smooth strongly monotone functions. One must add (7.24), which links both properties, and ensures the existence of some M -Lipschitz continuous function $h(x)$ interpolating h_i and h_j , and whose integral is smaller than or equal to $g_j - g_i$.

Class $\mathcal{S}_{M,+} = \int^{(2)} \mathcal{F}_{M,\frac{3}{2},+}$ of self-concordant functions

Corollary 7.3. A set $S = \{(x_i, g_i, h_i)\}_{i \in [N]} \in (\mathbb{R} \times \mathbb{R} \times \mathbb{R} \times \mathbb{R})^N$ is $\mathcal{S}_{M,+}$ -interpolable if and only if, $\forall i, j \in [N]$, $h_i = 0$ and $g_i = g_j$, or $\forall i, j \in [N]$,

$$|\tilde{h}_j - \tilde{h}_i| \leq M|x_j - x_i| \text{ and } h_i > 0 \tag{7.25}$$

If $\tilde{h}_i + \tilde{h}_j > -M(x_j - x_i)$, then

$$g_j - g_i \geq \frac{1}{M\tilde{h}_i} + \frac{1}{M\tilde{h}_j} - \frac{4}{M(\tilde{h}_i + \tilde{h}_j + M(x_j - x_i))}, \tag{7.26}$$

where $\tilde{h}_i = h_i^{-1/2}$.

Proof. Consider Theorem 7.3 with $\alpha = \frac{3}{2}$. \square

Remark 7.6. The class of M -self-concordant functions is affine-invariant, namely, if $f(x) \in \mathcal{S}_{M,+}$ then $f(ax + b) \in \mathcal{S}_{M,+} \forall a, b$ [NN94]. The conditions of Corollary 7.3 satisfy this affine-invariant property. Indeed, if a set $S = \{(x_i, g_i, h_i)\}$ is $\mathcal{S}_{M,+}$ -interpolable, then so is the set $\{(\frac{x_i - b}{a}, ag_i, a^2 h_i)\}$.

Class $\mathcal{T}_{M,+} = \int^{(2)} \mathcal{F}_{M,1,+}$ of quasi-self-concordant functions

Corollary 7.4. A set $S = \{(x_i, g_i, h_i)\}_{i \in [N]} \in (\mathbb{R} \times \mathbb{R} \times \mathbb{R} \times \mathbb{R})^N$ is $\mathcal{T}_{M,+}$ -interpolable if and only if, $\forall i, j \in [N]$, $h_i \geq 0$ and

$$g_j - g_i \geq \frac{h_i + h_j}{M} - \frac{2}{M} \sqrt{h_i h_j} e^{-\frac{M}{2}(x_j - x_i)}. \tag{7.27}$$

Proof. Consider Theorem 7.2. \square

Remark 7.7. The class of M -quasi-self-concordant is scale-invariant, namely, if $f \in \mathcal{T}_{M,+}$ then $cf \in \mathcal{T}_{M,+} \forall c > 0$. The conditions of Corollary 7.3 satisfy this affine-invariant property. Indeed, if a set $S = \{(x_i, g_i, h_i)\}$ is $\mathcal{S}_{M,+}$ -interpolable, then so is the set $\{(x_i, cg_i, ch_i)\}$.

It is known [Doi23, Lemma 2.7] that quasi-self-concordant functions satisfy

$$f'(y) - f'(x) - f''(x)(y - x) \leq \frac{1}{M} f''(x) (e^{M|y-x|} - M|y - x| - 1). \quad (7.28)$$

We rely on (7.27) to strengthen this condition:

Lemma 7.1. *If $f \in \mathcal{T}_{M,+}$, then $\forall x, y \in \mathbb{R}$,*

$$\begin{aligned} f'(y) - f'(x) - f''(x)(y - x) &\leq \frac{1}{M} f''(x) \left(e^{M|y-x|} - M|y - x| - 1 \right) \\ &\quad - \frac{1}{M} \left(\sqrt{f''(y)} - \sqrt{f''(x)e^{M(y-x)}} \right)^2. \end{aligned} \quad (7.29)$$

Proof. By Corollary 7.4, Remark 7.4, and using the identity $e^{|t|} - e^t - |t| + t \geq 0$, it holds that f satisfies, $\forall x, y \in \mathbb{R}$,

$$\begin{aligned} &f'(y) - f'(x) \\ &\leq -\frac{f''(x)}{M} - \frac{1}{M} \left(f''(y) - 2\sqrt{f''(x)f''(y)e^{M(y-x)}} \right) \\ &= -\frac{f''(x)}{M} - \frac{1}{M} \left(-f''(x)e^{M(y-x)} + \left(\sqrt{f''(y)} - \sqrt{f''(x)e^{M(y-x)}} \right)^2 \right) \\ &\leq -\frac{f''(x)}{M} - \frac{1}{M} \left(-f''(x)e^{M|y-x|} + f''(x)M(|y - x| - (y - x)) \right. \\ &\quad \left. + \left(\sqrt{f''(y)} - \sqrt{f''(x)e^{M(y-x)}} \right)^2 \right). \end{aligned}$$

□

7.3.4 Interpolation conditions for \mathcal{H}_M and $\mathcal{H}_{M,0}$.

Following the approach of Chapter 6, we obtain interpolation conditions with function values for \mathcal{H}_M and $\mathcal{H}_{M,0}$. Figure 7.4 illustrates the case \mathcal{H}_M .




Fig. 7.4 Derivation of interpolation conditions ensuring \mathcal{H}_M -interpolability of $S = \{(x_i, f_i, g_i, h_i)\}_{i=1,2}$. The middle figure displays the extremal M -smooth interpolants g_{\min} (dark blue) and g_{\max} (pink) of $\{(x_i, g_i, h_i)\}_{i=1,2}$. The top figure displays their derivatives, interpolating $\{(x_i, h_i)\}_{i=1,2}$. The bottom figure displays the integrated extremal interpolants (dark blue and pink), which define an interval including all possible functions in \mathcal{H}_M interpolating S , except possibly f_2 . The interval $[f_1 + \int_{x_1}^{x_2} g_{\min}(x) dx, f_1 + \int_{x_1}^{x_2} g_{\max}(x) dx]$ (light blue line, lower figure) is exactly the interval of admissible values for f_2 . Whenever f_2 belongs to this interval, there exists a convex combination of the integrals of g_{\min} and g_{\max} (dotted curve) which is in \mathcal{H}_M and interpolates S .

Since the proofs of Theorems 7.4 and 7.5, presenting interpolation conditions for \mathcal{H}_M and $\mathcal{H}_{M,0}$, respectively, follow the same steps as for the derivation of interpolation conditions for $\int \mathcal{F}_{M,\alpha,+}$, we defer them to Section 7.A. We directly present the interpolation conditions for \mathcal{H}_M and $\mathcal{H}_{M,0}$, and compare them with the existing literature.

Interpolation conditions for \mathcal{H}_M

Theorem 7.4 (Interpolation conditions for \mathcal{H}_M)

A set $S = \{(x_i, f_i, g_i, h_i)\}_{i \in [N]} \in (\mathbb{R} \times \mathbb{R} \times \mathbb{R} \times \mathbb{R})^{N+1}$ is \mathcal{H}_M -interpolable, if $\forall i, j \in [N]$

$$|h_j - h_i| \leq M|x_j - x_i| \quad (7.30)$$

If $h_j - h_i + M|x_j - x_i| \neq 0$, then

$$\begin{aligned} f_j - f_i - g_i(x_j - x_i) - \frac{h_i}{2}(x_j - x_i)^2 & \quad (7.31) \\ \geq -\frac{M}{6}|x_j - x_i|^3 & \\ & + \frac{\left(g_j - g_i - h_i(x_j - x_i) + \frac{M}{2}|x_j - x_i|(x_j - x_i)\right)^2}{2(h_j - h_i + M|x_j - x_i|)} \\ & + \frac{(h_j - h_i + M|x_j - x_i|)^3}{96M^2}, \end{aligned}$$

If $h_j - h_i + M|x_j - x_i| = 0$, then

$$g_j - g_i - h_i(x_j - x_i) = -\frac{M}{2}|x_j - x_i|(x_j - x_i), \quad (7.32)$$

$$f_j - f_i - g_i(x_j - x_i) - \frac{h_i}{2}(x_j - x_i)^2 = -\frac{M}{6}|x_j - x_i|^3. \quad (7.33)$$

Remark 7.8. Condition (7.31) is consistent with several multivariate conditions, that might serve as multivariate candidate interpolation conditions. However, we have, to this day, no insight into which candidate could be an interpolation condition, nor into how to prove so.

As already stated in $(p_{\nabla^2, M})$, it is known [Nes18, Lemma 1.2.4] that functions in \mathcal{H}_M satisfy a cubic bound, i.e., $\forall x, y \in \mathbb{R}$,

$$|f(y) - f(x) - f'(x)(y - x) - \frac{f''(x)}{2}(y - x)^2| \leq \frac{M}{6}|y - x|^3. \quad (7.34)$$

We show that the conditions of Theorem 7.4 strengthen (7.34). In addition, we show that a set satisfying these conditions necessarily satisfies Condition (7.21), as required by Theorem 6.1, and a lower bound on $f_j - f_i$, symmetrical to (7.31).

Lemma 7.2. *Consider a set $S = \{(x_i, g_i, h_i, f_i)\}_{i \in [N]}$. If S satisfies (7.30) and (7.31) for both pairs, then for both pairs S satisfies (7.21); and*

$$\begin{aligned} f_j - f_i - g_i(x_j - x_i) - \frac{h_i}{2}(x_j - x_i)^2 & \quad (7.35) \\ & \leq \frac{M}{6}|x_j - x_i|^3 \\ & \quad - \frac{(g_j - g_i - h_i(x_j - x_i) - M|x_j - x_i|(x_j - x_i))^2}{2(M|x_j - x_i| - h_j - h_i)} \\ & \quad - \frac{(M|x_j - x_i| - (h_j - h_i))^3}{96M^2}. \end{aligned}$$

In addition, (7.35) and (7.31) strengthen (7.34).

Proof. We shorten the notation of the pair $((x_i, g_i, h_i, f_i), (x_j, g_j, h_j, f_j))$ to (i, j) and let $\Delta x_{ij} = x_j - x_i$, $\Delta h_{ij} = h_j - h_i$, and

$$\begin{aligned} T_{ij}^g &= g_j - g_i - h_i(x_j - x_i), \\ T_{ij}^f &= f_j - f_i - g_i(x_j - x_i) - \frac{h_i}{2}(x_j - x_i)^2. \end{aligned}$$

Then, it holds

$$T_{ij}^g = -T_{ji}^g + \Delta h_{ij} \Delta x_{ij}, \quad (7.36)$$

$$T_{ij}^f = -T_{ji}^f + T_{ij}^g \Delta x_{ij} - \frac{1}{2} \Delta h_{ij} \Delta x_{ij}^2. \quad (7.37)$$

(Satisfaction of (7.35)) Condition (7.31) evaluated at (j, i) is equivalent to

Condition (7.35) evaluated at (i, j) :

$$\begin{aligned}
 T_{ji}^f &\geq -\frac{M}{6}|\Delta x_{ji}|^3 + \frac{\left(T_{ji}^g + \frac{M}{2}\Delta x_{ji}|\Delta x_{ji}|\right)^2}{2(\Delta h_{ji} + M|\Delta x_{ji}|)} + \frac{(\Delta h_{ji} + M|\Delta x_{ji}|)^3}{96M^2} \\
 \stackrel{(7.37)}{\Leftrightarrow} -T_{ij}^f + T_{ij}^g \Delta x_{ij} - \frac{1}{2}\Delta h_{ij}\Delta x_{ij}^2 &\geq -\frac{M}{6}|\Delta x_{ij}|^3 \\
 \stackrel{(7.36)}{\Leftrightarrow} T_{ij}^f &\leq \frac{M}{6}|\Delta x_{ij}|^3 - \frac{\left(T_{ij}^g - \frac{M}{2}\Delta x_{ij}|\Delta x_{ij}|\right)^2}{2(M|\Delta x_{ij}| - \Delta h_{ij})} - \frac{(M|\Delta x_{ij}| - \Delta h_{ij})^3}{96M^2} \\
 &\quad + \frac{\left(-T_{ij}^g + \Delta h_{ij}\Delta x_{ij} - \frac{M}{2}\Delta x_{ij}|\Delta x_{ij}|\right)^2}{2(-\Delta h_{ij} + M|\Delta x_{ij}|)} \\
 &\quad + \frac{(-\Delta h_{ij} + M|\Delta x_{ij}|)^3}{96M^2}
 \end{aligned}$$

Hence, if S satisfies (7.31) $\forall i, j = 1, 2$, then it satisfies (7.35) $\forall i, j = 1, 2$.

(Satisfaction of (7.21)) We now show that satisfaction $\forall i, j = 1, 2$ of (7.30) and (7.31) (hence satisfaction of (7.35)) imply satisfaction of (7.21) $\forall i, j = 1, 2$. Observe that (7.21) is satisfied by both pairs (i, j) and (j, i) if and only if

$$T_{ij}^g - K \geq 0 \text{ and } T_{ji}^g - K \geq 0 \stackrel{(7.36)}{\Leftrightarrow} -T_{ij}^g + \Delta h_{ij}\Delta x_{ij} - K \geq 0, \quad (7.38)$$

where $K = -\frac{M}{2}\Delta x_{ij}^2 + \frac{1}{4M}(\Delta h_{ij} + M\Delta x_{ij})^2$. Condition (7.38) is equivalent to

$$(T_{ij}^g - K)(-T_{ij}^g + \Delta h_{ij}\Delta x_{ij} - K) \geq 0, \quad (7.39)$$

where the equivalence follows from the fact that $T_{ij}^g < K$ and $-T_{ij}^g + \Delta h_{ij}\Delta x_{ij} < K$ is in contradiction with Condition (7.30), since it implies:

$$2K > \Delta h_{ij}\Delta x_{ij} \Leftrightarrow -\frac{M}{2}\Delta x_{ij}^2 + \frac{1}{2M}\Delta h_{ij}^2 > 0.$$

Subtracting (7.31) and (7.35) yields:

$$\frac{M}{3}|\Delta x_{ij}|^3 \geq \frac{\left(T_{ij}^g - \frac{M}{2}\Delta x_{ij}|\Delta x_{ij}|\right)^2}{2(M|\Delta x_{ij}| - \Delta h_{ij})} + \frac{\left(T_{ij}^g + \frac{M}{2}\Delta x_{ij}|\Delta x_{ij}|\right)^2}{2(\Delta h_{ij} + M|\Delta x_{ij}|)}$$

$$\begin{aligned}
 & + \frac{(M|\Delta x_{ij}| - \Delta h_{ij})^3 + (\Delta h_{ij} + M|\Delta x_{ij}|)^3}{96M^2} \\
 \Leftrightarrow & \frac{M|\Delta x_{ij}|}{(M|\Delta x_{ij}|)^2 - \Delta h_{ij}^2} \\
 & \left(-(T_{ij}^g)^2 + T_{ij}^g \Delta h_{ij} \Delta x_{ij} - \frac{M^2 \Delta x_{ij}^2}{16} - \frac{3 \Delta h_{ij}^2 \Delta x_{ij}^2}{8} + \frac{\Delta h_{ij}^4}{16M^2} \right) \geq 0 \\
 \stackrel{(7.30)}{\Leftrightarrow} & \left(T_{ij}^g - \frac{\Delta h_{ij}^2}{4M} + \frac{M \Delta x_{ij}^2}{4} - \frac{\Delta h_{ij} \Delta x_{ij}}{2} \right) \\
 & \left(-T_{ij}^g - \frac{\Delta h_{ij}^2}{4M} + \frac{M \Delta x_{ij}^2}{4} + \frac{\Delta h_{ij} \Delta x_{ij}}{2} \right) \geq 0 \\
 \Leftrightarrow & (T_{ij}^g - K)(-T_{ij}^g + \Delta h_{ij} \Delta x_{ij} - K) \geq 0,
 \end{aligned}$$

and (7.21) is satisfied.

(Strengthening of (7.34)) Finally, since $\Delta h_{ij} \geq -M|\Delta x_{ij}|$, (7.31) implies $T_{ij}^f \geq -\frac{M}{6}|\Delta x_{ij}|^3$, and since $\Delta h_{ij} \leq M|\Delta x_{ij}|$, (7.35) implies $T_{ij}^f \leq \frac{M}{6}|\Delta x_{ij}|^3$. \square

Remark 7.9. Conditions (7.21) and (7.35) can be removed from the final interpolation conditions since they are implied by the two remaining conditions. On the contrary, (7.30) cannot be removed from the interpolation conditions since it is not implied by (7.31). Indeed, consider for example the set $S = \{(0, 0, 0, 0), (\frac{1}{4}, \frac{1}{3}, 0, 0)\}$ satisfying (7.31) but not (7.30) for $M = 1$.

Condition (7.31) is not straightforward to manipulate. We thus propose a relaxation of this condition, satisfied by all functions in \mathcal{H}_M .

Corollary 7.5. *If f is a univariate twice-differentiable function in \mathcal{H}_M , then $\forall x, y \in \mathbb{R}$ such that $|f''(x) - f''(y)| \neq M|x - y|$, we have:*

$$\begin{aligned}
 & f(y) - f(x) - f'(x)(y - x) - \frac{f''(x)}{2}(y - x)^2 \tag{7.40} \\
 \leq & \frac{M}{6}|y - x|^3 - \frac{M}{3} \left(\frac{|f'(y) - f'(x) - f''(x)(y - x) - M|y - x|(y - x)|}{M} \right)^{\frac{3}{2}}.
 \end{aligned}$$

Proof. By Remark 7.4 and Lemma 7.2, $f \in \mathcal{H}_M$ satisfies (7.35) everywhere. In addition, for any $a, b, c \geq 0$ it holds:

$$3\frac{a}{c} + \frac{c^3}{b} \geq 4 \left(\frac{a^3 c^3}{c^3 b} \right)^{1/4} = 4 \left(\frac{a^3}{b} \right)^{1/4}. \tag{7.41}$$

Letting $a = \frac{1}{6}(f'(y) - f'(x) - f''(x)(y - x) - M|y - x|(y - x))^2$, $b = 96M^2$, and $c = M|y - x| - (f''(x) - f''(y))$ yields (7.40). \square

Inequality (7.40) will be exploited in Section 7.4, in Theorem 7.7.

Interpolation conditions for $\mathcal{H}_{M,\mu}$ For the sake of clarity, we consider the case $\mu = 0$. Interpolation conditions for $\mathcal{H}_{M,\mu}$, $\mu > 0$ can be recovered by considering $\tilde{f}(x) = f(x) - \frac{\mu}{2}x^2$.

Theorem 7.5 (Interpolation conditions for $\mathcal{H}_{M,0}$)

Given a set $S = \{(x_i, f_i, g_i, h_i)\}_{i \in [N]}$, let $\Delta x_{ij} = x_j - x_i$, $\Delta h_{ij} = h_j - h_i$, $T_{ij}^g = g_j - g_i - h_i(x_j - x_i)$, and $T_{ij}^f = f_j - f_i - g_i(x_j - x_i) - \frac{h_i}{2}(x_j - x_i)^2$. Then, S is $\mathcal{H}_{M,0}$ -interpolable if and only if $\forall i, j \in [N]$

$$h_i \geq 0 \tag{7.42}$$

$$\text{If } \Delta x_{ij} \geq \frac{h_i + h_j}{M}, \text{ then } g_j - g_i \geq \frac{h_i^2 + h_j^2}{2M}, \tag{7.43}$$

and,

1) If $\Delta h_{ij} + M|\Delta x_{ij}| \neq 0$, then

$$|\Delta h_{ij}| \leq M|\Delta x_{ij}| \tag{7.44}$$

$$T_{ij}^f \geq -\frac{M}{6}|\Delta x_{ij}|^3 + \frac{\left(T_{ij}^g + \frac{M}{2}\Delta x_{ij}|\Delta x_{ij}|\right)^2}{2(\Delta h_{ij} + M|\Delta x_{ij}|)} + \frac{(\Delta h_{ij} + M|\Delta x_{ij}|)^3}{96M^2} \tag{7.45}$$

If (7.48), then

$$f_j - f_i - g_i \Delta x_{ij} \geq \left(\frac{h_j^2 - h_i^2}{2M^2} + \frac{g_j - g_i}{M}\right) \left(M\sqrt{\frac{h_j^2 - h_i^2}{2M^2} + \frac{g_j - g_i}{M}} - h_j\right) + \frac{h_j^3 - h_i^3}{6M^2} + \frac{h_i^2 \Delta x_{ij}}{2M} \tag{7.46}$$

If (7.49), then

$$f_j - f_i - g_j \Delta x_{ij} \leq -\left(\frac{h_i^2 - h_j^2}{2M^2} + \frac{g_j - g_i}{M}\right) \left(M\sqrt{\frac{h_i^2 - h_j^2}{2M^2} + \frac{g_j - g_i}{M}} - h_i\right)$$

$$+ \frac{h_j^3 - h_i^3}{6M^2} - \frac{h_j^2 \Delta x_{ij}}{2M}, \quad (7.47)$$

where (7.48) and (7.49) are given by, respectively:

$$T_{ij}^g < \frac{(\Delta h_{ij} + M\Delta x_{ij})^2}{4M} - \frac{\Delta h_{ij}^2}{2M} - \frac{h_i}{M}(\Delta h_{ij} + M\Delta x_{ij}) \quad (7.48)$$

$$T_{ij}^g < \frac{(\Delta h_{ij} + M\Delta x_{ij})^2}{4M} - \frac{\Delta h_{ij}^2}{2M} - \frac{h_j}{M}(-\Delta h_{ij} + M\Delta x_{ij}). \quad (7.49)$$

2) If $\Delta h_{ij} + M|\Delta x_{ij}| = 0$, then

$$T_{ij}^g = -\frac{M}{2} \Delta x_{ij} |\Delta x_{ij}| \quad (7.50)$$

$$T_{ij}^f = -\frac{M}{6} |\Delta x_{ij}|^3. \quad (7.51)$$

Conditions (7.42) and (7.43) are exactly the interpolation conditions (7.23) and (7.24) ensuring interpolation of $\{(x_i, g_i, h_i)\}$ by an increasing M -smooth function. Then, (7.44) and (7.45) are exactly the interpolation conditions (7.30) and (7.31) for the class \mathcal{H}_M , and hence ensure consistency of S with a function whose Hessian is Lipschitz. Finally, (7.46) and (7.47) strengthen (7.45) (if needed), to ensure convexity of f .

7.4 Performance estimation of second-order optimization

This section uses the interpolation conditions derived in Section 7.3, i.e., Theorems 7.2 and 7.3, their Corollaries 7.1, 7.2, 7.3, and 7.4, and Theorems 7.4 and 7.5; to analytically and numerically analyze second-order methods. We aim at providing a global picture of the possible applications of our approach, rather than exhaustively analyzing all possible settings.

In particular, we show that *analytical* proofs of convergence can be obtained by

- Replacing, in existing convergence proofs, classical conditions by interpolation conditions, see, e.g., Theorem 7.7 for an analysis of (CNM) on \mathcal{H}_M . The convergence rates are better than existing ones.

- Solving PEPs by hand, see, e.g., Theorem 7.9 for an analysis of (GNM1) on $\mathcal{T}_{M,+}$. The convergence rates are *tight*.
- Identifying, from the numerical PEP, the worst-case instance, when such worst-case instance matches a known analytical upper bound, see, e.g., Theorem 7.6 proving tightness of an existing convergence rate of (NM) on \mathcal{H}_M ; Lemma 7.3 proving tightness of Theorem 7.9 for a single iteration; and Theorem 7.8 proving tightness of an existing convergence rate of (DNM) on $\mathcal{S}_{M,+}$. Whenever an upper bound holds in the multivariate setting, its tightness is also preserved in that setting.

In addition, solving PEPs (with interpolation conditions) numerically, without recovering the associated analytical proofs, yields

- *Numerical* exact convergence rates, see, e.g., Figures 7.5 and 7.6.
- Lower bounds for the multivariate setting, associated to the worst-case instances achieving the worst-case performance, see, e.g., Figure 7.5.
- An automated approach to compare methods on a fair basis, i.e., in the same setting, see, e.g., Figure 7.9.

Numerically, we use the solver Gurobi 11 [Gur23] to solve PEP formulations as in (7.4). This solver handles all interpolation conditions of Section 7.3, as constraints in the associated PEP formulations. In addition, it allows considering a broad range of performance measures, initial conditions, and methods, far beyond the scope of Gram-representable quantities. Computational constraints, on the other hand, often restrict our numerical analysis to one or two iterations.

7.4.1 Second-order methods on \mathcal{H}_M

Newton’s method: proof of tightness of existing rate. It is well-known that Newton’s method achieves a local quadratic convergence on \mathcal{H}_M , as stated, e.g., in [Nes18].

Proposition 7.5 (Theorem 1.2.5 of [Nes18]). *If*

- $f \in \mathcal{H}_M$,
- $\exists x_\star$ such that $\nabla f(x_\star) = 0$, $\nabla^2 f(x_\star) = \mu I \succ 0$,

- $\frac{M}{\mu} \|x_0 - x_\star\| \leq \frac{2}{3}$,

then Newton's iterations $x_{k+1} = x_k - \nabla^2 f(x_k)^{-1} \nabla f(x_k)$ satisfy

$$\|x_{k+1} - x_\star\| \leq \frac{\frac{M}{\mu} \|x_k - x_\star\|^2}{2 \left(1 - \frac{M}{\mu} \|x_k - x_\star\|\right)}. \quad (7.52)$$

To the best of our knowledge, the tightness of Proposition 7.5 is an open question. Relying on a PEP analysis, we identify a simple univariate function attaining (7.52), establishing its tightness for any number of iterations.

Theorem 7.6 (Tightness of [Nes18], Theorem 1.2.5)

Proposition 7.5 (i.e., [Nes18, Theorem 1.2.5]) is tight.

Proof. Consider

$$f(x) = -M \frac{|x|^3}{6} + \mu \frac{x^2}{2}.$$

Then, $x_\star = 0$ and $x_{k+1} = \frac{-\frac{M}{\mu} x_k |x_k|}{2(1 - \frac{M}{\mu} |x_k|)}$ attain (7.52) when $\frac{M}{\mu} |x_k| < 1$. \square

Such surprisingly simple worst-case functions were already observed, for example, for the gradient descent applied to smooth (strongly) convex functions [THG17c].

Cubic Newton's method: an improved descent lemma In the univariate case, CNM's iteration is given by

$$x_{k+1} = \underset{x}{\text{Arg min}} f(x_k) + f'(x_k)(x - x_k) + \frac{1}{2} f''(x_k)(x - x_k)^2 + \frac{M}{6} |x - x_k|^3, \quad (\text{CNM})$$

i.e., it consists in minimizing the quadratic approximation of the function at the current iterate, regularized by a cubic term. On \mathcal{H}_M , this ensures the monotonicity of CNM, i.e., [NP06, (2.11), Lemma 5]

$$f(x_{k+1}) \leq f(x_k) - \frac{M}{12} (\mathcal{P}(x_{k+1}))^3, \quad (7.53)$$

where

$$\mathcal{P}(x) = \max \left\{ \sqrt{\frac{|f'(x)|}{M}}, \frac{2}{3} f''(x) \right\}, \quad (7.54)$$

is the classical performance measure for CNM, as introduced in [NP06]. The resulting classical convergence rate resulting from the iterative application of (7.53), is reminded in Proposition 7.6.

Proposition 7.6 ([NP06], Theorem 1). *Given $N \in \mathbb{N}$, let $(x_k)_{k \in [N]}$ be the iterates of (CNM), with respect to some function $f \in \mathcal{H}_M$, and let f be bounded below by f_* . Then,*

$$\min_{k \in [N]} \mathcal{P}(x_k) \leq 2 \left(\frac{3(f(x_0) - f_*)}{2NM} \right)^{\frac{1}{3}}. \quad (7.55)$$

Using the refined description \mathcal{H}_M derived in Theorem 7.4, and in particular the improved cubic bound of Corollary 7.5, allows improving the descent lemma (7.53) for univariate functions by a factor of 5. This yields an improvement of factor $5^{\frac{1}{3}}$ on the global convergence rate (7.55) of (CNM) on univariate functions.

Theorem 7.7 (Improved convergence rate for CNM on \mathcal{H}_M)

Given $N \in \mathbb{N}$, let $(x_k)_{k=0,1,\dots,N}$ be the iterates of (CNM) on some function $f \in \mathcal{H}_M$. Then,

$$f(x_k) - f(x_{k+1}) \geq \frac{5M}{12} (\mathcal{P}(x_{k+1}))^3. \quad (7.56)$$

Moreover, if the function is bounded below by f_ , then*

$$\min_{k=0,1,\dots,N} \mathcal{P}(x_k) \leq \frac{2}{5^{\frac{1}{3}}} \left(\frac{3(f(x_0) - f_*)}{2NM} \right)^{\frac{1}{3}}. \quad (7.57)$$

Proof. The CNM iterates satisfy ([NP06, Equation (2.5), Proposition 1, and Lemmas 2 and 3])

$$f'(x_k) + f''(x_k)(x_{k+1} - x_k) + \frac{M}{2}(x_{k+1} - x_k)|x_{k+1} - x_k| = 0, \quad (7.58)$$

$$f''(x_k) + \frac{M}{2}|x_{k+1} - x_k| \geq 0, \quad (7.59)$$

$$f'(x_k)(x_{k+1} - x_k) \leq 0, \quad (7.60)$$

$$M|x_{k+1} - x_k|^2 \geq |f'(x_{k+1})|. \quad (7.61)$$

Claim 1: $\mathcal{P}(x_{k+1}) = \sqrt{\frac{|f'(x_{k+1})|}{M}} \Leftrightarrow \sqrt{\frac{|f'(x_{k+1})|}{M}} \geq -\frac{2}{3M}f''(x_{k+1})$.

Suppose first that $|x_{k+1} - x_k| = x_{k+1} - x_k \geq 0$. We have

$$\begin{aligned} |f'(x_{k+1})| &\geq -f'(x_{k+1}) \\ &\stackrel{(7.21)}{\geq} -f'(x_k) - \frac{M(x_{k+1} - x_k)^2}{4} + \frac{(f''(x_{k+1}) - f''(x_k))^2}{4M} \\ &\quad - \frac{1}{2}(f''(x_{k+1}) + f''(x_k))(x_{k+1} - x_k) \\ &\stackrel{(7.58)}{=} \frac{M(x_{k+1} - x_k)^2}{4} + \frac{(f''(x_{k+1}) - f''(x_k))^2}{4M} \\ &\quad - \frac{1}{2}(f''(x_{k+1}) - f''(x_k))(x_{k+1} - x_k) \\ &= \frac{1}{4M}(f''(x_{k+1}) - f''(x_k) - M(x_{k+1} - x_k))^2 \end{aligned}$$

Hence,

$$\begin{aligned} \sqrt{\frac{|f'(x_{k+1})|}{M}} &\geq \frac{|f''(x_{k+1}) - f''(x_k) - M|x_{k+1} - x_k||}{2M} \\ &\geq \frac{(M|x_{k+1} - x_k| - f''(x_{k+1}) + f''(x_k))}{2M} \\ &= \frac{\left(\frac{2}{3}M|x_{k+1} - x_k| + \frac{1}{3}M|x_{k+1} - x_k| - f''(x_{k+1}) + f''(x_k)\right)}{2M} \\ &\stackrel{(7.59)}{\geq} \frac{\left(-\frac{4}{3}f''(x_k) + \frac{4}{3}(f''(x_k) - f''(x_{k+1}))\right)}{2M} - \frac{2}{3M}f''(x_{k+1}). \end{aligned}$$

The case $|x_{k+1} - x_k| = -(x_{k+1} - x_k) \geq 0$ follows from a similar argument.

Claim 2: $f(x_k) - f(x_{k+1}) \geq \frac{5M}{12} \sqrt{\frac{|f'(x_{k+1})|}{M}}^3$. Exploiting the improved cubic

bound (7.40) of Corollary 7.5 yields

$$\begin{aligned}
 f(x_{k+1}) &\stackrel{(7.40)}{\leq} f(x_k) + f'(x_k)(x_{k+1} - x_k) + \frac{f''(x_k)}{2}(x_{k+1} - x_k)^2 \\
 &\quad + \frac{M}{6}|x_{k+1} - x_k|^3 - \frac{M^{\frac{1}{3}}}{3} \left(-|f'(x_{k+1}) - f'(x_k)| \right. \\
 &\quad \quad \left. - f''(x_k)(x_{k+1} - x_k) - \frac{M}{2}(x_{k+1} - x_k)|x_{k+1} - x_k| \right)^{\frac{3}{2}} \\
 &\stackrel{(7.58)}{=} f(x_k) + \frac{f'(x_k)}{2}(x_{k+1} - x_k) - \frac{M}{12}|x_{k+1} - x_k|^3 \\
 &\quad - \frac{M}{3} \left(\frac{|f'(x_{k+1})|}{M} \right)^{\frac{3}{2}} \\
 &\stackrel{(7.60)}{\leq} f(x_k) - \frac{M}{12}|x_{k+1} - x_k|^3 - \frac{M}{3} \left(\frac{|f'(x_{k+1})|}{M} \right)^{\frac{3}{2}} \\
 &\stackrel{(7.61)}{\leq} f(x_k) - \frac{M}{12} \left(\frac{|f'(x_{k+1})|}{M} \right)^{\frac{3}{2}} - \frac{M}{3} \left(\frac{|f'(x_{k+1})|}{M} \right)^{\frac{3}{2}} \\
 &= f(x_k) - \frac{5M}{12} \left(\frac{|f'(x_{k+1})|}{M} \right)^{\frac{3}{2}},
 \end{aligned}$$

establishing the result. Finally, telescoping the new descent lemma (7.56) yields the convergence rate in (7.57)

$$\begin{aligned}
 f(x_0) - f_* &\geq f(x_0) - f(x_N) = \sum_{k=0}^{N-1} f(x_k) - f(x_{k+1}) \\
 &\stackrel{(7.56)}{\geq} \frac{5M}{12} \sum_{k=0}^{N-1} (\mathcal{P}(x_{k+1}))^3 \geq \frac{5MN}{12} \min_{k=0, \dots, N-1} (\mathcal{P}(x_{k+1}))^3.
 \end{aligned}$$

□

Discussion on the tightness of Theorem 7.7

The improved descent lemma (7.56), and the associated convergence rate (7.57) for a single iteration, are tight.

Lemma 7.3. *Theorem 7.7 is tight for a single iteration.*

Proof. Let $f(x) = M\frac{x^3}{6} - \frac{x^2}{2}$ and $x_0 = 0$. Then, (CNM) yields

$$x_1 = \operatorname{Arg\,min}_x \frac{1}{2}(-x^2 + \frac{M}{3}|x|^3) = \pm \frac{2}{M}.$$

Let $x_1 = -\frac{2}{M}$. Then, $f(x_1) = -\frac{10}{3M^2}$, $f'(x_1) = \frac{4}{M}$ and

$$f(x_0) - f(x_1) = \frac{5M}{12} \left(\frac{|f'(x_1)|}{M} \right)^{\frac{3}{2}} = \frac{10}{3M^2}.$$

□

However, this tightness is lost when considering several iterations of (CNM). Figure 7.5 compares the worst-case guarantees on CNM, as computed by PEP and as predicted by [NP06, Theorem 1] and by Theorem 7.7.

Lower bound on the general worst-case performance

Solving a PEP with interpolation conditions, in the univariate setting, yields tight numerical convergence rates, but also an univariate worst-case function that achieves these rates. This worst-case function provides a lower bound on the general worst-case performance. Computing such lower bounds is an important direction of research, and the state-of-the-art lower bound for CNM was presented in [CGT10], via the exhibition of a univariate function (green dots in Figure 7.5). The lower bound resulting from solving the PEP (blue dots in Figure 7.5) is worse than that of [CGT10], and hence a better lower bound, obtained in an automated way.

7.4.2 Second order methods on $\mathcal{T}_{M,+}$

We now analyze the convergence of Newton's method (NM) and damped Newton methods (DNM) on standard self-concordant functions, i.e., $M = 1$. The classical performance measure in this setting is the *Newton decrement* [NN94], given by

$$\rho_f(x_k) = \frac{|f'(x_k)|}{\sqrt{f''(x_k)}} \quad (7.62)$$

Since neither the methods nor the performance criterion involve function values, the interpolation conditions of Theorem 7.3 can be used as such in PEP formulations, which yield exact univariate results. To consider, as in, e.g., [IH24], a performance criterion involving function values, and obtain

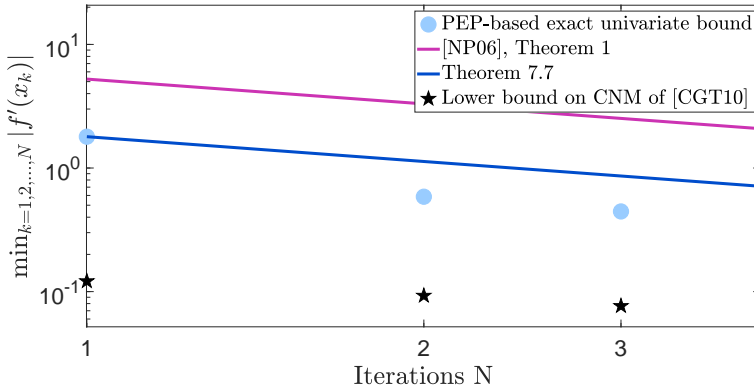


Fig. 7.5 Comparison of worst-case performance guarantees of (CNM) on \mathcal{H}_M , for $M = 1$. The performance measure is $\min_{k=0,1,\dots,N} |f'(x_k)|$ and the initial condition $f(x_0) - f_* \leq 1$. Plotted are the bound from [NP06, Theorem 1], the bound from Theorem 7.7, the exact univariate PEP-based worst-case bound, which serves as a multivariate lower bound, and an existing lower bound from [CGT10].

an a priori tight guarantee on this problem, one should derive interpolation conditions with function values for $\mathcal{T}_{M,+}$.

Both (NM) and (DNM) have been tightly analyzed for this performance measure, and for a single iteration, in [Hil21]. However, we use our approach to analyze the potential tightness of this single-iteration rate, on several iterations.

Newton's method On (NM), the exact single iteration convergence rate has been derived in [Hil21, Equation (11)].

Proposition 7.7 ([Hil21], Equation (11)). *Let $f \in \mathcal{S}_{M,+}$ with $M = 1$, and $\rho_f(x_k) \leq 1$. Then, $\rho_f(x_{k+1})$, where x_{k+1} is given by (NM), satisfies*

$$\rho_f(x_{k+1}) \leq 4 - \rho_f(x_k)^2 - 4\sqrt{1 - \rho_f(x_k)^2}. \quad (7.63)$$

Moreover, this bound is tight.

Solving the associated PEP allows identifying a worst-case function at-

taining the bound (7.7) for a single iteration, i.e., starting from $x_k = 0$,

$$f(x) = \begin{cases} Bx - \log(x - A) & \text{if } x \leq \frac{1}{2}(A - R) \\ \left(B + \frac{4}{R+A}\right)x - \log(x + R) & \text{if } x > \frac{1}{2}(A - R), \end{cases} \quad (7.64)$$

where $R = \rho_f(x_k)$, $A = R \frac{-R^2 + (-4 + 2\sqrt{1-R^2})R + 5 - \sqrt{1-R^2}}{R^2 - 1 + 2\sqrt{1-R^2}}$, and $B = \frac{R^2 - 1 + 2\sqrt{1-R^2}}{R(R-1)}$.

Worst-case performance for multiple iterations

As for Theorem 7.7 for (CNM), while Theorem 7.7 is tight for a single iteration, it can be improved for multiple iterations. Figure 7.6 displays an analysis of two iterations of (NM) on $\mathcal{S}_{M,+}$.




Fig. 7.6 Comparison of worst-case performance guarantees of two iterations of (NM) on $\mathcal{S}_{M,+}$. The performance measure $\rho_f(x_2)$ is plotted as a function of the initial condition $\rho_f(x_0)$. Displayed are the exact PEP-based bound (blue dots), and the bound of [Hil21, Equation (11)] (pink).

It compares PEP-based results, and the exact single iteration bound 7.7 from [Hil21, Equation (11)]. The exact performance is better than predicted by both theoretical bounds, and the region of convergence is larger.

Damped Newton method On (DNM), the exact single iteration convergence rate has also been derived in [Hil21, Equation (11)].

Proposition 7.8 ([Hil21], Equation (11)). *Let $f \in \mathcal{S}_{M,+}$ with $M = 1$, and $\alpha \leq$*

$2 \frac{\sqrt{1+\rho_f(x_k)^3}-1}{\rho_f(x_k)^3}$. Then, $\rho_f(x_{k+1})$, where x_{k+1} is given by (DNM), satisfies

$$\rho_f(x_{k+1}) \leq \rho_f(x_k) - \alpha \rho_f(x_k) + \alpha \rho_f(x_k)^2. \quad (7.65)$$

Moreover, this bound is tight.

Again, solving the associated PEP allows identifying a worst-case function attaining the bound (7.8) for a single iteration, i.e., starting from $x_k = 0$,

$$f(x) = \frac{R-1}{R}x - \log(R-x) \quad (7.66)$$

where $R = \rho_f(x_k)$.

Moreover, we show (7.65) is actually tight even for multiple iterations.

Theorem 7.8 (Convergence of (DNM) on $\mathcal{S}_{M,+}$, $N \geq 2$)

Let $f \in \mathcal{S}_{M,+}$ with $M = 1$. Given $N \in \mathbb{N}$, let $\alpha_k \leq 2 \frac{\sqrt{1+\rho_f(x_k)^3}-1}{\rho_f(x_k)^3}$, $k = 0, 1, \dots, N$, and let $(x_k)_{k=0,1,\dots,N}$ be the iterates of (DNM) with respect to f and α_k . Then,

$$\rho_f(x_{k+1}) \leq \rho_f(x_k) - \alpha_k \rho_f(x_k) + \alpha_k \rho_f(x_k)^2 \quad \forall k \in [N-1]. \quad (7.67)$$

Moreover, this bound is tight.

Proof. The bound follows directly from Proposition 7.7, and (7.66) achieves this bound. \square

7.4.3 Second-order methods on $\mathcal{T}_{M,+}$

Gradient Regularized Newton Method

GNM1 [Doi23], see also [Pol09], is a variant of Newton's method with the following explicit iteration

$$x_{k+1} = x_k - \frac{f'(x_k)}{f''(x_k) + M|f'(x_k)|}. \quad (\text{GNM1})$$

This method, on $\mathcal{T}_{M,+}$, achieves a linear rate of convergence [Doi23, Theorem 3.3], and exhibits a quadratic local convergence on quasi-self-concordant

functions [Doi23, Theorem 4.1]. The proof of this quadratic convergence relies on the following descent lemma.

Lemma 7.4 ([Doi23], Equation (49)). *Let $f \in \mathcal{T}_{M,+}$, $\eta(x) = M \frac{|f'(x)|}{f''(x)}$, and $(x_k)_{k=0,1,\dots}$ by given by (GNM1). Then,*

$$\eta(x_{k+1}) \leq e^{\eta(x_k)} (e^{\eta(x_k)} + \eta(x_k)^2 - \eta(x_k) - 1). \quad (7.68)$$

Analytically solving the associated PEP allows improving on this descent lemma. Figure 7.7 illustrates the difference between relying on exact interpolation conditions and necessary conditions when solving a PEP.

Theorem 7.9 (Descent Lemma for (GNM1) on $\mathcal{T}_{M,+}$)

Let $f \in \mathcal{T}_{M,+}$, $\eta(x) = M \frac{|f'(x)|}{f''(x)}$, and $(x_k)_{k=0,1,\dots}$ by given by (GNM1). Then,

$$\eta(x_{k+1}) \leq e^{\frac{\eta(x_k)}{\eta(x_k)+1}} (\eta(x_k) - 1) + 1 \quad (7.69)$$

where $\eta(x) = M \frac{|f'(x)|}{f''(x)}$.

Proof. The PEP associated to this problem is

$$\begin{aligned} \max_{S=\{(x_0, g_0, h_0), (x_1, g_1, h_1)\}} \quad & \frac{|g_1|}{h_1} \text{ s.t. } x_1 = x_0 - \frac{g_0}{h_0 + M|g_0|}, \\ & \frac{|g_0|}{h_0} = R, \\ & S \text{ is } \mathcal{T}_{M,+}\text{-interpolable,} \end{aligned} \quad (7.70)$$

Figure 7.7 illustrates the difference in the feasible domains of (7.70), when $\mathcal{T}_{M,+}$ -interpolability of S is enforced via classical conditions defining quasi-self-concordant functions, i.e., (7.28), or the interpolation conditions from Corollary 7.4.

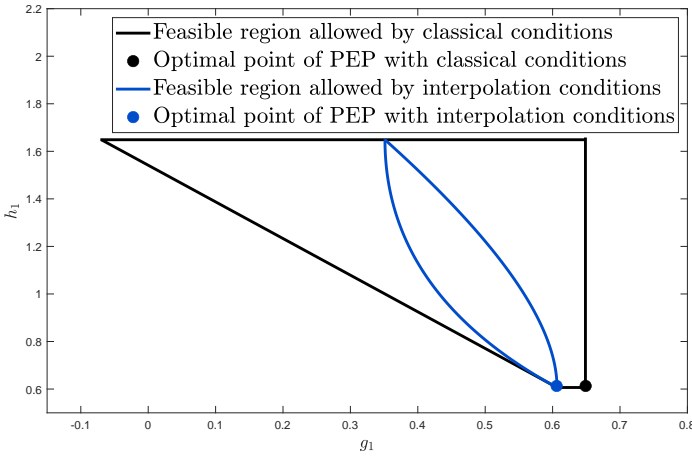


Fig. 7.7 Given $(x_0, g_0, h_0) = (0, 1, 1)$, $M = 1$ and $x_1 = \frac{1}{2}$, admissible region for (g_1, h_1) such that $\{(x_i, g_i, h_i)\}_{i=0,1}$ satisfies (i) classical conditions (7.28), (black curves), or (ii) interpolation conditions from Corollary 7.4 (red curves). It also shows the points (g_1, h_1) corresponding to the worst-case of (GNM1), according to both conditions. Using interpolation conditions instead of necessary conditions significantly restricts the domain (g_1, h_1) of the PEP, and the optimal point obtained by solving a PEP with necessary conditions (black dot) is not consistent with any actual function of the class.

W.l.o.g., assume $x_0 = 0$ and $g_0 = 1$, since this problem is invariant under translation, and scaling of f . The two equality constraints in (7.70) yield $h_0 = \frac{1}{R}$, and $x_1 = \frac{-R}{MR+1}$. The interpolation conditions are (see Theorem 7.4)

$$g_1 - g_0 \geq \frac{h_0 + h_1}{M} - \frac{2}{M} \sqrt{h_0 h_1} e^{-\frac{M}{2}(x_1 - x_0)}$$

$$g_0 - g_1 \geq \frac{h_1 + h_0}{M} - \frac{2}{M} \sqrt{h_1 h_0} e^{-\frac{M}{2}(x_0 - x_1)}.$$

One can check that the optimal solution to (7.70), satisfying these conditions exactly, is

$$h_1^* = h_0 e^{-\frac{\eta_0}{\eta_0+1}}, \quad g_1^* = g_0 + \frac{-h_0 + h_1^*}{M},$$

yielding the optimal value

$$M \frac{|g_1^*|}{h_1^*} = e^{\frac{\eta_0}{\eta_0+1}} (\eta_0 - 1) + 1,$$

where $\eta_0 = MR = M \frac{|g_0|}{h_0}$. □

Figure 7.8 compares this PEP-based bound with [Doi23, Equation (49)].




Fig. 7.8 Worst-case performance of one iteration of (GNM1) on $\mathcal{T}_{M,+}$, $M = 1$. The performance measure $\eta(x_1) = M \frac{|f'(x_1)|}{f''(x_k)}$ is plotted as a function of the initial condition $\eta(x_0) = M \frac{|f'(x_0)|}{f''(x_0)}$. We compare the PEP-based numerical solutions (blue dots), the known descent lemma of [Doi23, Equation (49)] (pink curve), and Theorem 7.9 (blue curve).

7.4.4 Second-order methods on $\mathcal{H}_{M,\mu}$

Comparison of methods on $\mathcal{H}_{M,\mu}$ Second-order methods are usually analyzed on settings specific to the method, i.e., on different function classes, under different initial conditions, and for different performance measures. This makes it difficult to compare methods fairly. By contrast, a strength of the PEP framework is that it allows straightforwardly obtaining convergence rates in any setting, and hence renders such comparisons possible. As a simple example, Figure 7.9 compares the worst-case performance of

a single iteration of several second-order methods on $\mathcal{H}_{M,\mu}$, with performance measure $|f'(x_1)|$ and initial condition $|f'(x_k)| \leq 1$. Specifically, it compares Newton's method (NM), Cubic Newton's method (CNM), and Gradient regularized method (GNM1). In this setting, (CNM) achieves the fastest convergence rate. In addition, (NM) achieves quadratic local convergence for μ sufficiently large only, while (CNM) and (GNM1) converge independently of μ .




Fig. 7.9 Comparison of worst-case performance of an iteration of different second-order methods, i.e., (NM), (CNM), and (GNM1), on $\mathcal{H}_{M,\mu}$ with μ varying and $M = 1$. The performance measure is $|f'(x_1)|$ and the initial condition $|f'(x_k)| \leq 1$.

7.5 Concluding observations

This chapter took a step forward into tightly analyzing second-order optimization methods, by enabling tight analysis of the *univariate* setting, via (i) the derivation of interpolation conditions for a range of second-order function classes, and (ii) the resolution of the associated non-convex PEP formulations, either analytically or numerically.

On the downside, the approach suffers mainly from (i) its restriction to the univariate setting, (ii) the somewhat tedious derivation of interpolation conditions, especially when considering function values, and (iii) the computational burden of solving the associated PEP formulations, hence the restriction to the analysis of a very few number of iterations.

Appendix to Chapter 7

7.A Proof of Theorems 7.4 and 7.5

For the sake of readability, we let $\Delta x_{ij} = x_j - x_i$, $\Delta x_{xi} = x - x_i$, $\Delta h_{ij} = h_j - h_i$, $T_{ij}^g = g_j - g_i - h_i(x_j - x_i)$, and $T_{ij}^f = f_j - f_i - g_i(x_j - x_i) - \frac{h_i}{2}(x_j - x_i)^2$, and use

$$T_{ij}^g = -T_{ji}^g + \Delta h_{ij} \Delta x_{ij}, \quad (7.71)$$

$$T_{ij}^f = -T_{ji}^f + T_{ij}^g \Delta x_{ij} - \frac{1}{2} \Delta h_{ij} \Delta x_{ij}^2. \quad (7.72)$$

By Proposition 7.2, \mathcal{H}_M and $\mathcal{H}_{M,0}$ are convex and order 2 connectable. Building on the interpolation conditions for $\int \mathcal{F}_M$, and $\int \mathcal{F}_{M,\mu}$, given in Corollaries 7.1 and 7.2, we now show that $\int \mathcal{F}_M$ and $\int \mathcal{F}_{M,0}$ are extremally interpolable, and provide their extremal interpolants.

Lemma 7.5 (Extremal interpolants of \mathcal{H}_M). *Consider an interval $[x_1, x_2]$ and a \mathcal{H}_M -interpolable pair of points $S = \{(x_i, f_i, g_i, h_i)\}_{i=1,2}$. Then, in $[x_1, x_2]$, the expressions of g_{\min} and g_{\max} as defined in (6.7) are given by:*

When $\Delta h_{12} \neq \pm M \Delta x_{12}$,

- [g_{\min} , Case 1] If we consider $\int \mathcal{F}_M$, or if $y_1 \leq x_1 + \frac{h_1}{M}$,

$$g_{\min}(x) = \begin{cases} g_1 + h_1 \Delta x_{x1} - \frac{M}{2} \Delta x_{x1}^2 & \text{if } x \in [x_1, y_1] \\ g_1 + h_1 \Delta x_{x1} - \frac{M}{2} (y_1 - x_1)^2 \\ \quad + \frac{M}{2} (x^2 - y_1^2) + M(x_1 - 2y_1)(x - y_1) & \text{if } x \in [y_1, y_2] \\ g_2 + h_2 \Delta x_{x2} - \frac{M}{2} \Delta x_{x2}^2 & \text{if } x \in [y_2, x_2], \end{cases} \quad (7.73)$$

where

$$y_1 = x_2 - \frac{T_{12}^g + \frac{M}{2}\Delta x_{12}^2}{\Delta h_{12} + M\Delta x_{12}} - \frac{\Delta h_{12} + M\Delta x_{12}}{4M}$$

$$y_2 = x_2 - \frac{T_{12}^g + \frac{M}{2}\Delta x_{12}^2}{\Delta h_{12} + M\Delta x_{12}} + \frac{\Delta h_{12} + M\Delta x_{12}}{4M}.$$

• [g_{\min} , Case 2] Else,

$$g_{\min}(x) = \begin{cases} g_1 + h_1\Delta x_{x_1} - \frac{M}{2}\Delta x_{x_1}^2 & \text{if } x \in [x_1, x_1 + \frac{h_1}{M}] \\ g_1 + h_1(w_1 - x_1) - \frac{M}{2}(w_1 - x_1)^2 & \text{if } x \in [x_1 + \frac{h_1}{M}, w_1] \\ g_1 + h_1(w_1 - x_1) - \frac{M}{2}(w_1 - x_1)^2 \\ \quad + \frac{M}{2}(x^2 - w_1^2) - M(w_1x - w_1^2) & \text{if } x \in [w_1, w_2] \\ g_2 + h_2\Delta x_{x_2} - \frac{M}{2}\Delta x_{x_2}^2 & \text{if } x \in [w_2, x_2], \end{cases} \quad (7.74)$$

where

$$w_1 = x_2 + \frac{h_2}{M} - 2\sqrt{\frac{g_2 - g_1}{M} + \frac{h_2^2 - h_1^2}{2M^2}}$$

$$w_2 = x_2 + \frac{h_2}{M} - \sqrt{\frac{g_2 - g_1}{M} + \frac{h_2^2 - h_1^2}{2M^2}}.$$

• [g_{\max} , Case 1] If we consider $\int \mathcal{F}_M$, or if $y_2 \geq x_2 - \frac{h_2}{M}$,

$$g_{\max}(x) = \begin{cases} g_1 + h_1\Delta x_{x_1} + \frac{M}{2}\Delta x_{x_1}^2 & \text{if } x \in [x_1, z_1] \\ g_1 + h_1(z_1 - x_1) + \frac{M}{2}(z_1 - x_1)^2 \\ \quad - \frac{M}{2}(x^2 - z_1^2) - M(x_1 - 2z_1)(x - z_1) & \text{if } x \in [z_1, z_2] \\ g_2 + h_2\Delta x_{x_2} + \frac{M}{2}\Delta x_{x_2}^2 & \text{if } x \in [z_2, x_2], \end{cases} \quad (7.75)$$

where

$$z_1 = x_2 + \frac{T_{12}^g - \frac{M}{2}\Delta x_{12}^2}{-\Delta h_{12} + M\Delta x_{12}} - \frac{(-\Delta h_{12} + M\Delta x_{12})}{4M}$$

$$z_2 = x_2 + \frac{T_{12}^g - \frac{M}{2}\Delta x_{12}^2}{-\Delta h_{12} + M\Delta x_{12}} + \frac{(-\Delta h_{12} + M\Delta x_{12})}{4M}.$$

• [g_{\max} , Case 2] Else,

$$g_{\max}(x) = \begin{cases} g_1 + h_1 \Delta x_{x_1} + \frac{M}{2} \Delta x_{x_1}^2 & \text{if } x \in [x_1, v_1] \\ g_1 + h_1(v_1 - x_1) + \frac{M}{2}(v_1 - x_1)^2 \\ -\frac{M}{2}(x^2 - v_1^2) + M(xv_1 - v_1^2) & \text{if } x \in [v_1, v_2] \\ g_2 + h_2(x - v_2) + \frac{M}{2}(x - v_2)^2 & \text{if } x \in [v_2, x_2 - \frac{h_2}{M}] \\ g_2 + h_2 \Delta x_{x_2} + \frac{M}{2} \Delta x_{x_2}^2 & \text{if } x \in [x_2 - \frac{h_2}{M}, x_2], \end{cases} \quad (7.76)$$

where

$$v_1 = x_1 - \frac{h_1}{M} + \sqrt{\frac{g_2 - g_1}{M} - \frac{h_2^2 - h_1^2}{2M^2}}$$

$$v_2 = x_1 - \frac{h_1}{M} + 2\sqrt{\frac{g_2 - g_1}{M} - \frac{h_2^2 - h_1^2}{2M^2}}.$$

When $\Delta h_{12} = \pm M \Delta x_{12}$:

$$g_{\max}(x) = g_{\min}(x) = g_1 + h_1 \Delta x_{x_1} \pm \frac{M}{2} \Delta x_{x_1}^2 = g_2 + h_2 \Delta x_{x_2} \pm \frac{M}{2} \Delta x_{x_2}^2. \quad (7.77)$$

Proof. Suppose $\Delta h_{12} \neq \pm M \Delta x_{12}$, and consider (7.73). For any $x \in [x_1, x_2]$, let

$$g'_{\min}(x) = \begin{cases} h_1 - M \Delta x_{x_1} & \text{if } x \in [x_1, y_1] \\ h_1 + M(x + x_1 - 2y_1) & \text{if } x \in [y_1, y_2] \\ h_2 - M \Delta x_{x_2} & \text{if } x \in [y_2, x_2], \end{cases} \quad (7.78)$$

Then, $\forall x \in [x_1, x_2]$, $g_{\min}(x) = \int_{x_1}^x g'_{\min}(t) dt$. If $x \in [x_1, y_2]$, this is straightforward, and if $x \in [y_2, x_2]$ it follows from:

$$\begin{aligned} \int_{x_1}^x g'_{\min}(t) dt &= g_1 + h_1(y_2 - x_1) - \frac{M}{2}(y_1 - x_1)^2 + \frac{M}{2}(y_2^2 - y_1^2) \\ &\quad + M(x_1 - 2y_1)(y_2 - y_1) + h_2(x - y_2) \\ &\quad - \frac{M}{2}(x^2 - y_2^2) + Mx_2(x - y_2) \\ &= g_2 + h_2 \Delta x_{x_2} - \frac{M}{2} \Delta x_{x_2}^2 - T_{12}^g - \Delta h_{12}(y_2 - x_2) \end{aligned}$$

$$\begin{aligned}
 & -M \left((y_2 - x_2)\Delta x_{12} + \frac{\Delta x_{12}^2}{2} - (y_2 - y_1)^2 \right) \\
 &= g_2 + h_2\Delta x_{x2} - \frac{M}{2}\Delta x_{x2}^2 - (T_{12}^g + M\frac{\Delta x_{12}^2}{2}) \\
 & \quad - (\Delta h_{12} + M\Delta x_{12})(y_2 - x_2) + \frac{\Delta h_{12} + M\Delta x_{12}}{4M} \\
 &= g_2 + h_2\Delta x_{x2} - \frac{M}{2}\Delta x_{x2}^2.
 \end{aligned}$$

We first observe that $g_{\min}(x)$ interpolates S since

$$g_{\min}(x_1) = g_1, \quad g_{\min}(x_2) = g_2, \quad g'_{\min}(x_1) = h_1 \quad \text{and} \quad g'_{\min}(x_2) = h_2.$$

In addition, $g_{\min} \in \int \mathcal{F}_M$, since $g'_{\min} \in \mathcal{F}_M$. Indeed, (i) $g'_{\min} \in \mathcal{C}^0$ and is piecewise \mathcal{C}^1 , and (ii) it is a linear function by part, of slope M . Inclusion in \mathcal{F}_M follows from order 0-connectability of \mathcal{F}_M .

It remains to show that $g_{\min}(x)$ and $g'_{\min}(x)$ are the minimal g and associated h satisfying (7.21) with respect to S , which completes the proof. To this end, consider the following optimization problem: $g_{\min}(x) =$

$$\left\{ \begin{array}{l} \min_{g,h} \quad g \\ \text{s.t.} \quad g_i - g \geq \frac{1}{4M}(h - h_i)^2 - \frac{M}{4}\Delta x_{xi}^2 - \frac{1}{2}(h + h_i)\Delta x_{xi}, \quad i = 1,2 \quad (\mu_i) \\ \quad \quad g - g_i \geq \frac{1}{4M}(h_i - h)^2 - \frac{M}{4}\Delta x_{xi}^2 + \frac{1}{2}(h_i + h)\Delta x_{xi}, \quad i = 1,2 \quad (\lambda_i) \end{array} \right. \quad (7.79)$$

where μ_i, λ_i are the associated dual variables. We consider its KKT conditions.

$$\text{Stat. cond. in } g: 1 + \sum_i (\mu_i - \lambda_i) = 0$$

$$\begin{aligned}
 \text{Stat. cond. in } h: \quad & \sum_i \left(\mu_i \left(\frac{1}{2M}(h - h_i) - \frac{1}{2}\Delta x_{xi} \right) \right. \\
 & \left. + \lambda_i \left(\frac{1}{2M}(h - h_i) + \frac{1}{2}\Delta x_{xi} \right) \right) = 0
 \end{aligned}$$

$$\text{Dual feas.:} \quad \mu_i, \lambda_i \geq 0$$

$$\text{Primal feas.:} \quad g_i - g \geq \frac{1}{4M}(h - h_i)^2 - \frac{M}{4}\Delta x_{xi}^2 - \frac{1}{2}(h + h_i)\Delta x_{xi}, \quad i = 1,2 \quad (7.80)$$

$$g - g_i \geq \frac{1}{4M}(h_i - h)^2 - \frac{M}{4}\Delta x_{xi}^2 + \frac{1}{2}(h_i + h)\Delta x_{xi}, \quad i = 1, 2 \quad (7.81)$$

$$\begin{aligned} \text{Slackness:} \quad \mu_i(g_i - g - (\frac{1}{4M}(h - h_i)^2 - \frac{M}{4}\Delta x_{xi}^2 - \frac{1}{2}(h + h_i)\Delta x_{xi})) &= 0 \\ \lambda_i(g - g_i - (\frac{1}{4M}(h_i - h)^2 - \frac{M}{4}\Delta x_{xi}^2 + \frac{1}{2}(h_i + h)\Delta x_{xi})) &= 0. \end{aligned}$$

Let $x \in [x_1, y_1]$. Then, setting all dual coefficients zeros except for $\lambda_1 = 1$, $h = h_1 - M\Delta x_{x1}$ and $g = g_1 + h_1\Delta x_{x1} - \frac{M}{2}\Delta x_{x1}^2$ is a valid solution. Indeed, all conditions follow directly from the definition of g and h , except for (7.80), $i = 2$, which becomes equivalent to (7.21) and is thus satisfied by assumption on S , and (7.81), $i = 2$, which becomes:

$$\begin{aligned} T_{12}^g - \frac{3M}{4}\Delta x_{x1}^2 - \frac{M}{4}\Delta x_{x2} - \frac{M}{2}\Delta x_{x1}\Delta x_{x2} + \frac{\Delta h_{12}^2}{4M} + \frac{\Delta h_{12}}{2}(2x - x_1 - x_2) &\leq 0 \\ \Leftrightarrow \Delta x_{x2}(\Delta h_{12} + M\Delta x_{12}) &\leq -T_{12}^g - \frac{\Delta h_{12}^2}{4M} - \frac{3M\Delta x_{x1}^2}{4} - \frac{\Delta h_{12}\Delta x_{12}}{2} \\ \Leftrightarrow x &\leq y_1. \end{aligned}$$

Let $x \in [y_2, x_2]$. Then, setting all dual coefficients zeros except for $\lambda_2 = 1$, $h = h_2 - M\Delta x_{x2}$ and $g = g_2 + h_2\Delta x_{x2} - \frac{M}{2}\Delta x_{x2}^2$ is a valid solution. Again, all conditions follow directly from the definition of g and h , except for (7.80), $i = 1$, which becomes equivalent to (7.21) as applied to the pair (j, i) and is thus satisfied by assumption on S , and (7.81), $i = 1$, which becomes:

$$\begin{aligned} T_{12}^g - \frac{3M}{4}\Delta x_{x2}^2 + \frac{M}{4}\Delta x_{x1} + \frac{M}{2}\Delta x_{x1}\Delta x_{x2} - \frac{\Delta h_{12}^2}{4M} + \frac{\Delta h_{12}}{2}(2x + \frac{x_1}{2} - \frac{3x_2}{2}) &\geq 0 \\ \Leftrightarrow \Delta x_{x2}(\Delta h_{12} + M\Delta x_{12}) &\geq -T_{12}^g + \frac{\Delta h_{12}^2}{4M} - \frac{M\Delta x_{x1}^2}{4} + \frac{\Delta h_{12}\Delta x_{12}}{2} \\ \Leftrightarrow x &\geq y_2. \end{aligned}$$

Else, setting $\mu_1 = \mu_2 = 0$,

$$\begin{aligned} \lambda_2 &= \frac{2M(x - y_1)}{\Delta h_{12} + M\Delta x_{12}} \Big|_{x \geq y_1, (7.30)} \geq 0, \\ \lambda_1 &= 1 - \lambda_0 \Big|_{x \leq y_2} \geq 1 - \frac{2M(y_2 - y_1)}{\Delta h_{12} + M\Delta x_{12}} \stackrel{\text{Def. of } y_2}{=} 0 \\ h &= h_1 + M(x + x_1 - 2y_1) \end{aligned}$$

and

$$\begin{aligned} g &= g_1 + \frac{M}{4}(x + x_1 - 2y_1)^2 - \frac{M}{4}\Delta x_{x_1}^2 + \frac{1}{2}(2h_1 + M(x + x_1 - 2y_1))\Delta x_{x_1} \\ &= g_1 + h_1\Delta x_{x_1} - \frac{M}{2}(y_1 - x_1)^2 + \frac{M}{2}(x^2 - y_1^2) + M(x_1 - 2y_1)(x - y_1), \end{aligned}$$

is a valid solution. Indeed, all conditions follow directly from the definition of g , h , λ_0 and λ_1 except for (7.80), $i = 1, 2$, which become equivalent to (7.30) and are thus satisfied by assumption on S .

Suppose now $\Delta h_{12} = -M\Delta x_{12}$, which implies $T_{12}^g = -\frac{M}{2}\Delta x_{12}^2$ since (7.21) applied to both pairs becomes:

$$\begin{aligned} T_{12}^g &\geq -\frac{M}{2}\Delta x_{12}^2 \\ T_{ji}^g &\geq -\frac{M}{2}\Delta x_{12}^2 \stackrel{(7.36)}{\Leftrightarrow} T_{12}^g \leq \Delta h_{12}\Delta x_{12} + \frac{M}{2}\Delta x_{12}^2 = -\frac{M}{2}\Delta x_{12}^2. \end{aligned}$$

In this case, let $g'_{\min}(x) = h_1 - M\Delta x_{x_1} = h_2 - M\Delta x_{x_2}$. Clearly, $g'_{\min}(x)$ is the first derivative of $g_{\min}(x)$, and this function interpolates S . Moreover, in this case, $\mu_0 = \mu_1 = \lambda_1 = 0$, $\lambda_0 = 1$,

$$\begin{aligned} g &= g_1 + h_1\Delta x_{x_1} - \frac{M}{2}\Delta x_{x_1}^2 \stackrel{T_{12}^g = -\frac{M}{2}\Delta x_{12}^2}{=} g_2 + h_2\Delta x_{x_2} - \frac{M}{2}\Delta x_{x_2}^2 \\ h &= h_1 - M\Delta x_{x_1} = h_2 - M\Delta x_{x_2}. \end{aligned}$$

are always a solution satisfying the KKT conditions associated with Problem (7.79). Indeed, all conditions follow directly from these two definitions of g and h .

The computation of g_{\min} (case 2) and $g_{\max}(x)$ (cases 1 and 2) follow the same argument, except with the associated derivatives

$$g'_{\min}(x) = \begin{cases} h_1 - M\Delta x_{x_1} & \text{if } x \in [x_1, x_1 + \frac{h_1}{M}] \\ 0 & \text{if } x \in [x_1 + \frac{h_1}{M}, w_1] \\ M(x - w_1) & \text{if } x \in [w_1, w_2] \\ h_2 - M\Delta x_{x_2} & \text{if } x \in [w_2, x_2], \end{cases}$$

$$g'_{\max}(x) = \begin{cases} h_1 + M\Delta x_{x1} & \text{if } x \in [x_1, z_1] \\ h_1 - M(x + x_1 - 2z_1) & \text{if } x \in [z_1, z_2] \\ h_2 + M\Delta x_{x2} & \text{if } x \in [z_2, x_2]. \end{cases}$$

$$g'_{\max}(x) = \begin{cases} h_1 + M\Delta x_{x1} & \text{if } x \in [x_1, V_1] \\ h_1 - M(x - v_2) & \text{if } x \in [v_1, v_2] \\ 0 & \text{if } x \in [v_2, x_2 - \frac{h_2}{M}] \\ h_2 + M\Delta x_{x2} & \text{if } x \in [x_2 - \frac{h_2}{M}, x_2]. \end{cases}'$$

where, when considering $\int \mathcal{F}_{M,+}$, g_{\min} (resp. g_{\max}) is set to 0 when becoming non-positive. \square

Building on Lemma 7.5, we proceed to the proof of Theorem 7.4. Indeed, in Proposition 7.9, we prove necessity and sufficiency of (7.30), (7.31), (7.35) and (7.21) to \mathcal{H}_M -interpolability of a set $S = \{(x_i, g_i, h_i, f_i)\}_{i=1,2}$. By Lemma 7.2, these conditions are equivalent to the conditions of Theorem 7.4. And, by Lemma 6.5, these interpolation conditions hold for any set $S = \{(x_i, f_i, g_i, h_i)\}_{i \in [N]}$. Therefore, Proposition 7.9 concludes the proof of Theorem 7.4. The proof of Theorem 7.5 follows the same lines, except that it handles the different cases of extremal interpolants derived in Lemma 7.5

Proposition 7.9. *A set $S = \{(x_i, g_i, h_i, f_i)\}_{i=1,2}$ is \mathcal{H}_M -interpolable if and only if, $\forall i, j = 1, 2$,*

1) *If $\Delta h_{ij} + M|\Delta x_{ij}| \neq 0$, then*

$$|\Delta h_{ij}| \leq M|\Delta x_{ij}|, \quad (7.82)$$

$$T_{ij}^f \geq -\frac{M}{6}|\Delta x_{ij}|^3 + \frac{\left(T_{ij}^g + \frac{M}{2}\Delta x_{ij}|\Delta x_{ij}\right)^2}{2(\Delta h_{ij} + M|\Delta x_{ij}|)} + \frac{(\Delta h_{ij} + M|\Delta x_{ij}|)^3}{96M^2}, \quad (7.83)$$

$$T_{ij}^f \leq \frac{M}{6}|\Delta x_{ij}|^3 - \frac{\left(T_{ij}^g - \frac{M}{2}\Delta x_{ij}|\Delta x_{ij}\right)^2}{2(M|\Delta x_{ij}| - \Delta h_{ij})} - \frac{(M|\Delta x_{ij}| - \Delta h_{ij})^3}{96M^2}, \quad (7.84)$$

$$T_{ij}^g \geq -\frac{M}{2}\Delta x_{ij}^2 + \frac{1}{4M}(\Delta h_{ij} + M\Delta x_{ij})^2, \quad (7.85)$$

2) If $\Delta h_{ij} \pm M|\Delta x_{ij}| = 0$, then

$$T_{ij}^g = \mp \frac{M}{2} \Delta x_{ij} |\Delta x_{ij}| \quad (7.86)$$

$$T_{ij}^f = \mp \frac{M}{6} |\Delta x_{ij}|^3 \quad (7.87)$$

Proof. Suppose $x_1 < x_2$. By Lemma 6.1, it holds that S is \mathcal{F}_M -interpolable if and only if it satisfies $|h_j - h_i| \leq M|x_j - x_i|$ (Proposition 7.1), (7.85) and

$$\int_{x_1}^{x_2} g_{\min}(x) dx \leq f_2 - f_1 \leq \int_{x_1}^{x_2} g_{\max}(x) dx, \quad (7.88)$$

where g_{\min} and g_{\max} are given in Lemma 7.5.

Consider first the case $\Delta h_{12} \neq \pm M|\Delta x_{12}|$. Integration of extremal interpolants gives:

$$T_{12}^f \geq -\frac{M}{6} \Delta x_{12}^3 + \frac{(T_{12}^g + \frac{M}{2} \Delta x_{12}^2)^2}{2(\Delta h_{12} + M\Delta x_{12})} + \frac{(\Delta h_{12} + M\Delta x_{12})^3}{96M^2} \quad (7.89)$$

$$T_{12}^f \leq \frac{M}{6} \Delta x_{12}^3 - \frac{(T_{12}^g - \frac{M}{2} \Delta x_{12}^2)^2}{2(M\Delta x_{12} - \Delta h_{12})} - \frac{(M\Delta x_{12} - \Delta h_{12})^3}{96M^2}. \quad (7.90)$$

Suppose now $x_2 < x_1$. Then, taking Lemma 7.5 with x_2 and x_1 inverted and following the same argument as in the case $x_1 < x_2$, we obtain the following inequalities to be satisfied:

$$T_{21}^f \geq -\frac{M}{6} \Delta x_{21}^3 + \frac{(T_{21}^g + \frac{M}{2} \Delta x_{21}^2)^2}{2(\Delta h_{21} + M\Delta x_{21})} + \frac{(\Delta h_{21} + M\Delta x_{21})^3}{96M^2} \quad (7.91)$$

$$T_{21}^f \leq \frac{M}{6} \Delta x_{21}^3 + \frac{(T_{21}^g - \frac{M}{2} \Delta x_{21}^2)^2}{2(\Delta h_{21} - M\Delta x_{21})} + \frac{(\Delta h_{21} - M\Delta x_{21})^3}{96M^2}. \quad (7.92)$$

By (7.71) and (7.72), Condition (7.91) becomes:

$$\begin{aligned} -T_{12}^f + T_{12}^g \Delta x_{12} - \frac{1}{2} \Delta h_{12} \Delta x_{12}^2 &\geq \frac{M}{6} \Delta x_{12}^3 - \frac{(\Delta h_{12} + M\Delta x_{12})^3}{96M^2} \\ &\quad + \frac{(-T_{12}^g + \Delta h_{12} \Delta x_{12} + \frac{M}{2} \Delta x_{12}^2)^2}{-2(\Delta h_{12} + M\Delta x_{12})} \\ \Leftrightarrow -T_{12}^f + T_{12}^g \Delta x_{12} - \frac{1}{2} \Delta h_{12} \Delta x_{12}^2 &\geq \frac{M}{6} \Delta x_{12}^3 - \frac{(T_{12}^g + \frac{M}{2} \Delta x_{12}^2)^2}{2(\Delta h_{12} + M\Delta x_{12})} \end{aligned}$$

$$\begin{aligned}
 & + T_{12}^g \Delta x_{12} - \frac{1}{2} \Delta h_{12} \Delta x_{12}^2 \\
 & - \frac{(\Delta h_{12} + M \Delta x_{12})^3}{96M^2} \\
 \Leftrightarrow T_{12}^f \leq & -\frac{M}{6} \Delta x_{12}^3 + \frac{(T_{12}^g + \frac{M}{2} \Delta x_{12}^2)^2}{2(\Delta h_{12} + M \Delta x_{12})} + \frac{(\Delta h_{12} + M \Delta x_{12})^3}{96M^2} \quad (7.93)
 \end{aligned}$$

Similarly, (7.92) becomes:

$$\begin{aligned}
 -T_{12}^f + T_{12}^g \Delta x_{12} - \frac{1}{2} \Delta h_{12} \Delta x_{12}^2 & \leq -\frac{M}{6} \Delta x_{12}^3 - \frac{(\Delta h_{12} - M \Delta x_{12})^3}{96M^2} \\
 + \frac{(-T_{12}^g + \Delta h_{12} \Delta x_{12} - \frac{M}{2} \Delta x_{12}^2)^2}{-2(\Delta h_{12} - M \Delta x_{12})} \\
 \Leftrightarrow T_{12}^f \geq & \frac{M}{6} \Delta x_{12}^3 + \frac{(T_{12}^g - \frac{M}{2} \Delta x_{12}^2)^2}{2(\Delta h_{12} - M \Delta x_{12})} + \frac{(\Delta h_{12} - M \Delta x_{12})^3}{96M^2}. \quad (7.94)
 \end{aligned}$$

Conditions (7.89) and (7.94) are equivalent to a single constraint, independently of the sign of Δx_{12} :

$$T_{12}^f \geq -\frac{M}{6} |\Delta x_{12}|^3 + \frac{\left(T_{12}^g + \frac{M}{2} \Delta x_{12} |\Delta x_{12}|\right)^2}{2(\Delta h_{12} + M |\Delta x_{12}|)} + \frac{(\Delta h_{12} + M |\Delta x_{12}|)^3}{96M^2}.$$

Similarly, Conditions (7.90) and (7.93) can be expressed as:

$$T_{12}^f \leq \frac{M}{6} |\Delta x_{12}|^3 - \frac{\left(T_{12}^g - \frac{M}{2} \Delta x_{12} |\Delta x_{12}|\right)^2}{2(M |\Delta x_{12}| - \Delta h_{12})} - \frac{(M |\Delta x_{12}| - \Delta h_{12})^3}{96M^2}.$$

We have obtained interpolation conditions to be satisfied by the pair (1,2) for S to be \mathcal{H}_M -interpolable. Extending the argument to the pair (2,1) allows concluding that a \mathcal{H}_M -interpolable set $S = \{(x_i, g_i, h_i, f_i)\}_{i=1,2}$ necessarily satisfies, for all $i, j = 1, 2$, Conditions (7.83) and (7.84).

Suppose now $x_2 > x_1$ and $\Delta h_{12} = \pm M \Delta x_{12}$, $T_{12}^g = \pm \frac{M}{2} \Delta x_{12}^2$. Then,

$$f_2 - f_1 = \int_{x_1}^{x_2} g_{\min}(x) dx \Leftrightarrow T_{12}^f = \pm \frac{M}{6} \Delta x_{12}^3 \quad (7.95)$$

Finally, suppose $x_1 < x_2$ and $\Delta h_{12} = \pm M \Delta x_{12} \Leftrightarrow \Delta h_{21} = \pm M \Delta x_{21}$, implying $T_{12}^g = \mp \frac{M}{2} \Delta x_{12}^2$. By Lemma 7.5 applied to the pair (2,1), it holds

that

$$T_{21}^f = \pm \frac{M}{6} \Delta x_{21}^3 \stackrel{(7.36),(7.37)}{\Leftrightarrow} T_{12}^f = \mp \Delta x_{12}^3. \quad (7.96)$$

Combining (7.95) and (7.96), it holds that $\forall i, j = 1, 2$ if $\Delta h_{ij} = \pm M |\Delta x_{ij}|$, then necessarily $T_{ij}^f = \pm \frac{M}{6} |\Delta x_{ij}|^3$ (and $T_{ij}^g = \pm \frac{M}{2} \Delta x_{ij} |\Delta x_{ij}|$) for S to be \mathcal{H}_M -interpolable. \square

PART III
**Interpolation conditions for
stochastic first-order
optimization**

8

Computer-aided analysis of stochastic first-order methods

This chapter results from a collaboration with Sébastien Colla, and benefited from fruitful discussions with Balasz Gerencsér.

This chapter proposes an extension of the PEP framework, initially introduced to analyze deterministic methods, to the analysis of first-order stochastic methods. That is, we consider, to solve unconstrained minimization problems

$$\min_{x \in \mathbb{R}^d} f(x),$$

where $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is a convex function, *stochastic* first-order black-box optimization methods, which rely on evaluations of f and *noisy* evaluations of its gradient to iteratively minimize f . Such methods include, e.g., the stochastic gradient method (SGD) whose iteration is given by

$$x_{k+1} = x_k - \alpha_k(g_k + \varepsilon_k), \quad g_k \in \partial f(x_k) \quad (\text{SGD})$$

where ε_k is a zero-mean random variable, and independent across queries.

Our goal, given a stochastic first-order method \mathcal{M} , a function class \mathcal{F} , and a noise model for ε_k with prescribed expectation and variance, is to obtain a worst-case guarantee on the expected performance of \mathcal{M} on \mathcal{F} . That is, a guarantee valid for all functions in \mathcal{F} and all distributions of ε_k satisfying the noise model. In particular, we say the bound is *tight* when attained by explicit functions and noise distributions within the considered classes. This chapter presents a computer-aided framework that computes a range of such worst-case guarantees, whose tightness depends on the complexity of the corresponding formulation.

The chapter is organized as follows. Section 8.1 motivates the analysis of stochastic first-order methods, reviews computer-aided advances in this area, and outlines the proposed framework. Section 8.2 briefly presents some PEP tools on which the framework builds. Section 8.3 then introduces the problem classes covered by the framework and the framework itself. Finally, Section 8.4 applies the framework to analyze the stochastic gradient method (SGD) on the class $\mathcal{F}_{\mu,L}$ of smooth (strongly) convex functions, under various noise models—including additive bounded noise, finite-sum settings with or without variance reduction, and block-coordinate schemes. This illustrates, both analytically and numerically, how the framework aligns with and complements the existing literature on stochastic first-order methods.

Notation Given $k \in \mathbb{N}$, we denote by \mathbf{e}_k the k^{th} unit vector. We let $\mathcal{D}^d := \{M \in \mathbb{R}^{d \times d} \mid M \text{ diagonal}\}$. In addition, given $M \in \mathbb{R}^{d \times d}$, we let $\|M\| = \max_{x \in \mathbb{R}^d: \|x\|=1} \|Mx\|$.

Given N random variables $\varepsilon_0, \dots, \varepsilon_{N-1}$, we denote by ε the random variable that aggregates all ε_k 's, whose distribution depends on those of ε_k 's, i.e. $\varepsilon = [\varepsilon_0, \dots, \varepsilon_{N-1}]$. We denote by $\mathbb{E}[\cdot]$ the expectation with respect to the randomness of the method under analysis, that is, the expectation taken over the joint distribution of all random variables involved in the method. We let $\{\mathcal{A}_k\}_{k=0,1,\dots,N}$ be the filtration of available information at iteration k , see, e.g., [Dur19].

8.1 Motivation

Stochastic first-order methods, as opposed to exact first-order methods, i.e., $\varepsilon_k = 0, \forall k$, naturally arise in a wide range of contexts where one has access

to (or prefers to compute), e.g., an unbiased estimator of the gradient of f rather than the gradient itself. Depending on these contexts, the distribution of ε_k satisfies different assumptions.

Additive noise on deterministic gradients The simplest setting arises when minimizing a deterministic function whose gradient evaluations are corrupted by additional zero-mean stochastic noise. This noise can be either *relative*, with variance depending on the gradient norm ($\mathbb{E}_{\varepsilon_k}[\|\varepsilon_k\|^2] \leq \sigma^2 \|\nabla f(x_k)\|^2, \forall k$) or *absolute*, with variance independent of the gradient norm ($\mathbb{E}_{\varepsilon_k}[\|\varepsilon_k\|^2] \leq \sigma^2, \forall k$).

Such noise may arise in practice from simulation or measurement-based functions involving inherent randomness or error, or from intentional perturbations used to ensure data privacy, as in differentially private algorithms, see, e.g., [BST14].

Stochastic optimization problems Another example is the minimization of a expected value, i.e.,

$$f(x) = \mathbb{E}_{\zeta \sim \mathcal{D}}[f_{\zeta}(x)], \quad (8.1)$$

where samples $\zeta \sim \mathcal{D}$ are available, but the distribution \mathcal{D} itself is unknown. In such cases, gradients of f are unavailable, contrary to gradients of f_{ζ} , which provide an unbiased estimator of ∇f . Hence, methods rely on noisy unbiased evaluations $\nabla f_{\zeta}(x_k) = \nabla f(x_k) + \varepsilon_k$.

Stochastic optimization problems are particularly relevant in statistical supervised learning. Indeed, letting \mathcal{D} be the unknown distribution linking inputs and outputs, ζ be an input-output sample drawn from \mathcal{D} , x a prediction model, and $f_{\zeta}(x)$ the loss incurred by x on ζ , minimizing f corresponds to minimizing the expected risk or loss of x ; see, e.g., [BCN18, Section 3.1].

Finite-sum minimization One of the most studied categories of problems relying on stochastic optimization consists in the minimization of a large sum of functions, i.e.,

$$f(x) = \sum_{i=1}^n f_i(x), \quad (8.2)$$

where n is large and f_i 's are differentiable convex functions. In contrast with the first two examples, exact computation of ∇f is possible in this case, but is computationally expensive. On the other hand, choosing uniformly at random an index $i \in \{1, 2, \dots, n\}$, furnishes $\nabla f_i(x)$, an unbiased estimator of $\nabla f(x)$, whose computation is approximately n times cheaper than the one of $\nabla f(x)$. As shown, e.g., in [BB07], it can therefore be advantageous to minimize (8.2) using stochastic methods that rely on $\nabla f_i(x_k) = \nabla f(x_k) + \varepsilon_k$, instead of $\nabla f(x_k)$. Alternative ways to construct estimators of ∇f in this setting, leading to different assumptions on ε_k , are discussed in Section 8.1.1.

Such minimization problems typically arise as approximations of (8.1) using n independent samples of ζ , and are then referred to as empirical risk minimization (ERM) or sample average approximation (SAA) problems; see, e.g., [BCN18, Section 3.1].

Randomized block-coordinate descent methods Finally, another example consists in minimizing large-dimensional differentiable functions that exhibit component-wise structure. In this setting, it is again possible to compute the exact gradient, but it is computationally advantageous to evaluate the gradient only along a single or a block of coordinates, chosen uniformly at random. In this case, methods sample $d\nabla_i f(x_k)\mathbf{e}_i = \nabla f(x_k) + \varepsilon_k$, where \mathbf{e}_i is the i -th unit vector, d is the dimension of x , $\nabla_i f(x_k)$ is the i^{th} component of $\nabla f(x_k)$, and $\mathbb{E}_{\varepsilon_k}[\varepsilon_k] = 0$. Such methods are known as *randomized coordinate descent* (RCD) methods, and can be generalized to gradient estimators along blocks of coordinates; see, e.g., [Nes12, BT13].

8.1.1 Related works

Stochastic first-order methods First-order stochastic methods, first conceptualized in [RM51], are the subject of extensive research, particularly given their applications in machine learning and large-scale optimization problems, see, e.g., [Bor08, BCN18, GG23, GLQ⁺19, Net19] and references therein. Analyses depend on (i) the method, (ii) the function class considered, and (iii) the noise model, that is, at minima, its variance and expectation, and, in some structural settings as the finite sum and the block-coordinate settings, additional information on its distribution. We provide a non-exhaustive summary of existing convergence rates for stochastic first-order methods on convex and strongly convex functions, under different noise settings. These rates should be compared to the deterministic

case ($\varepsilon_k = 0$), in which the simplest method (GD) achieves, with constant stepsizes, rates of $\mathcal{O}(1/k)$ in terms of the function value gap on convex functions, with k the number of iterations, and linear convergence in terms of the distance to the minimizer, on strongly convex smooth functions, see, e.g., [Nes18].

A widely analyzed model, though restrictive in its applications, is the additive bounded noise model ($\mathbb{E}[\|\varepsilon_k\|^2] \leq \sigma^2$), see, e.g., [Nem94, NY83, NJLS09]. Under this model, methods ranging from SGD, with or without iterates averaging [PJ92, Rup88], to methods involving momentum [Xia09, GL12, CDO18, CHVSL18, DGN13] have been designed and analyzed on smooth convex functions. In this setting, (SGD), with iterates averaging in the convex case, yields, with decreasing stepsizes, $\mathcal{O}(1/\sqrt{k})$ and $\mathcal{O}(1/k)$ rates for convex and strongly convex functions. In addition, (SGD) with small constant stepsizes converges to a neighborhood depending on σ^2 , at optimal $\mathcal{O}(1/k)$ and linear rates, see, e.g., [MB11, RSS11, Sti19].

Stochastic first-order methods have also been extensively analyzed in the finite-sum setting $f = \frac{1}{n} \sum_i f_i$. Analyses depend on the properties of the f_i 's, and on the way an estimator of $\nabla f(x_k)$ is constructed at each iteration. Classical assumptions include convexity and L_i -smoothness of each f_i , and bounded variance at the optimum, i.e., $\mathbb{E}[\|\nabla f_i(x_*)\|^2] = \sigma_*^2 \leq +\infty$, where x_* is a minimizer of f . Under these assumptions, (SGD) where gradient estimators sample $\nabla f_i(x_k)$ uniformly (or via importance sampling) has been extensively analyzed [BCN18, GLQ⁺19, MB11, NSW14, SGD21, VBS19, ZZ15]. As for the bounded additive noise model, (SGD) with decreasing stepsizes yield $\mathcal{O}(1/\sqrt{k})$ and $\mathcal{O}(1/k)$ rates for convex (with iterates averaging) and strongly convex functions, see, e.g., [GG23, Theorem 5.7, Corollary 5.9]. In addition, small constant stepsizes yield optimal $\mathcal{O}(1/k)$ and linear rates to a neighborhood depending on σ_*^2 , see, e.g., [GG23, Theorems 5.6, 5.8]. In particular, when $\sigma_*^2 = 0$ (e.g., in the finite-sum setting for overparameterized models, see, e.g., [MBB18, VBS19]), (SGD) with constant stepsizes achieves optimal $\mathcal{O}(1/k)$ and linear rates, as in the deterministic setting. In addition, if estimators of $\nabla f(x_k)$ are built using variance-reduction techniques [DBLJ14, GSB20, NLST17, RSB12, SLRB17, JZ13, SSZ13, XZ14], i.e., so that the variance of the noise on $\nabla f(x_k)$ naturally decreases iteratively, mitigating the effect of σ_*^2 , then (SGD) achieves linear convergence on strongly convex functions. Examples include SAGA [DBLJ14] and SVRG [JZ13, XZ14]. In addition to (SGD), accelerated methods, see, e.g., [AZ18, LMH15, VBS19, ZDS⁺19], and proximal point methods [TASV24] have also been analyzed in this setting.

Further, stochastic first-order methods have been analyzed under weak and strong growth conditions. Respectively, $\mathbb{E}_{\varepsilon_k}[\|\varepsilon_k\|^2] \leq A(f(x_k) - f(x_\star))$ [VBS19, FR15], and $\mathbb{E}_{\varepsilon_k}[\|\varepsilon_k\|^2] \leq B\|\nabla f(x_k)\|^2$ [SR13], where x_\star is a minimizer of f , $(x_k)_{k=0,1,\dots}$ are the iterates of the method, and $A, B \geq 0$. These conditions yield optimal rates for (SGD) on convex and strongly convex functions. Those same convergence rates have been proven in the block-coordinate setting, see, e.g., [BT13, Nes12, Wri15].

Finally, beyond smooth convex optimization, stochastic proximal methods have also been designed and analyzed on *non-smooth* functions, see, e.g., [DS09, RVV20], and stochastic first-order methods on *non-convex* functions have been, more recently, investigated under several assumptions, see, e.g., [AD19, DDG⁺22, DD19, DG19, DS20, GL13, KR20]. In particular, [GL13] showed, for (SGD) on L -smooth functions, an $\mathcal{O}(\frac{1}{k^4})$ convergence rate to stationary points, under the additive bounded noise model; [DS20] provided matching lower bounds in this setting, and [DD19, DG19] showed similar rate of convergence on weakly convex functions.

Unified analysis of stochastic first-order methods As highlighted in, e.g., [GHR20, TASV24, Sti19], analyses of first-order methods across many settings can be unified by analyzing problem classes encoding (i) the method, (ii) the function class considered, and (iii) the expectation and variance of the noise ε_k . Such analyses yield convergence rates valid for all distributions satisfying the noise model. In particular, [GHR20, Assumption 4.1] introduces the class of noise models for which there exist non-negative constants $A_1, A_2, C, D_1, D_2, \rho \in \mathbb{R}$ and a possibly random sequence (σ_k^2) such that

$$\begin{aligned} \mathbb{E}[\varepsilon_k | x_k] &= 0, \\ \mathbb{E}[\|\varepsilon_k\|^2 | x_k] &\leq A_1(f(x_k) - f(x_\star)) - \|\nabla f(x_k)\|^2 + C\sigma_k^2 + D_1, \\ \mathbb{E}[\sigma_{k+1}^2 | x_k] &\leq (1 - \rho)\sigma_k^2 + A_2(f(x_k) - f(x_\star)) + D_2, \end{aligned} \quad (8.3)$$

where x_\star is a minimizer of f and $(x_k)_{k=0,1,\dots}$ are the iterates of the method. On smooth strongly convex functions, Assumption (8.3) suffices for linear convergence of (SGD) to a neighborhood depending on C_1, C_2 [GHR20, Theorem 4.1], and this assumption is satisfied by most settings mentioned above, see [GHR20, Table 2].

For instance, the finite-sum setting, where each f_i is convex and L_i -smooth, and gradient estimators uniformly sample $\nabla f_i(x_k)$, satisfies (8.3) with $A = 4 \max_i L_i$ and $B = 2\sigma_\star^2$; see, e.g., [GG23, Lemma 4.20]. Similarly,

the additive bounded noise model on L -smooth functions satisfies (8.3) with $A = 2L$ and $B = \sigma^2$; see, e.g., [Sti19, Example 2]. These shared parameters explain the similarity in convergence rates across these noise models. Likewise, the strong growth condition, the additive relative noise model, and the block-coordinate setting can all be captured as a single noise model under this unified noise framework, as they share the same expectation and variance.

Computer-aided analysis of stochastic first-order methods As discussed in Chapter 3, computer-aided analyses were, at first, introduced in deterministic setting, but there exist extensions to the stochastic setting that we quickly summarize.

First, computer-aided analyses naturally extend to the analysis of *inexact* methods, where gradient errors are deterministic perturbations, as in [d'A08, DGN13, DGN14], see, e.g., [DKGT20, BTB22, Gan22, HSL21, THG17a, VG24, VG25a, VG25b] for PEP analyses of inexact methods. We refer to this approach as the *worst-scenario* approach. Indeed, conducting such analyses on settings in which perturbations are, in fact, stochastic rather than deterministic, yields convergence rates that hold for the worst realization of the noisy perturbations, rather than in expectation. Such rates are thus pessimistic in these settings.

Second, when the noise takes values in a finite set of cardinality n , e.g., in the finite-sum and the block-coordinate settings, stochastic methods can be analyzed *exactly* by enumerating all scenarios and considering the average-case performance, see, e.g., [TB19] for an introduction to the all-scenarios approach, [KHG23] for block-coordinate methods, [HSR17, CKPG25] for the finite-sum setting. We refer to this approach by the *all-scenarios* approach. The resulting numerical convergence rates are tight, and the associated closed-form analytical rates may remain valid for distributions with infinite-dimensional supports, as showcased, e.g., in [TB19, Appendix E]. On the downside, the size of the associated SDPs grows like n^N , limiting the scalability of this approach and restricting its use to the analysis of a small number of iterations only. Moreover, when rates depend on the support cardinality n , extracting closed-form expressions requires solving a sequence of increasingly large problems, whose solutions hopefully provide intuition for the analytical expression.

Then, [VSL21] studied the additive bounded noise model by (i) theoretically linking deterministic SDP solutions to robustness and convergence guarantees (see, e.g., [VSL21, Theorems 5.3, 6.2]), and (ii) leveraging these

SDPs to analyze or design methods.

Finally, [AdKZ23b] analyzed the RCD by incorporating expectations of stochastic quantities, e.g., $\mathbb{E}[f(x_k)]$, $\mathbb{E}[\|x_k\|^2]$, or $\mathbb{E}[\|\varepsilon_k\|^2]$, directly as variables in the associated SDPs, under *necessary* constraints meant to ensure the variables are indeed expectations. See also [DT14, Appendix B] for early introduction of this same idea, under the additive bounded noise model.

8.1.2 Introduction to the proposed framework

Inspired by the idea of [AdKZ23b, DS20] to embed expectations of stochastic quantities as variables to PEP formulations, and by the unified approach to stochastic performance analysis proposed, e.g., in [GHR20], we present a PEP-based approach for analyzing all stochastic zero-mean noises, and such that, for some constants $A_1, B_1, C_1, E_1, A_2, B_2, C_2, D_2, \rho \in \mathbb{R}$, and a possibly random sequence $(\sigma_k^2)_{k=0,1,\dots,N-1}$, it holds

$$\begin{aligned} \mathbb{E}[\|\varepsilon_k\|^2 | \mathcal{A}_k] &\leq A_1(f(x_k) - f(x_\star)) + B_1\|g_k\|^2 + C_1\|x_k - x_\star\|^2 + D_1 + E_1\sigma_k^2, \\ \mathbb{E}[\sigma_{k+1}^2 | \mathcal{A}_k] &\leq A_2(f(x_k) - f(x_\star)) + B_2\|g_k\|^2 + C_2\|x_k - x_\star\|^2 + D_2 \\ &\quad + (1 - \rho)\sigma_k^2, \end{aligned} \tag{8.4}$$

where $(x_k)_{k=0,1,\dots,N}$ is the sequence of iterates of a first-order method, $g_k \in \partial f(x_k)$, x_\star is a minimizer of f , and $\mathcal{A}_k := \sigma(x_0, \dots, x_k, \varepsilon_0, \dots, \varepsilon_{k-1}, \sigma_0^2, \dots, \sigma_k^2)$ is the σ -algebra of available information up to iteration k . This model includes, as special cases, all classical noise assumptions presented in Section 8.1.1.

More precisely, we propose a framework, fitting the PEP methodology, for automated analysis of first-order stochastic methods with fixed step-sizes across a wide range of function classes (specified in Assumption 8.1), under all zero-mean noise models satisfying (8.4) (Section 8.3). Given N a number of iterations, a method, a function class, and a noise model (expectation and variance), we propose a spectrum of SDPs of increasingly large size, ranging from a size $\propto N$ (Section 8.3.2) to a size $\propto 2^N$ (Section 8.3.5), whose solutions provide increasingly stronger convergence rates on the problem. In particular, the SDP of size $\propto 2^N$ yields tight performance guarantees, i.e., that are attained by a specific function and noise distribution. On the other side of the spectrum, the SDP of size $\propto N$ yields guarantees

that often happen to be, numerically, already tight (Section 8.4).

By construction, on *structural* noise models, e.g., on the finite-sum or the block-coordinate settings, the all-scenarios approach introduced in [TB19], which uses this structural information on the noise's distribution, yields stronger convergence rates than our framework. However, (i) these rates are valid for a smaller class of problems, preventing a unified analysis across different settings, (ii) the all-scenarios approach is restricted to finite-support noises, and can only improve on our framework on models for which information on the noise distribution is available, and (iii) this approach is mainly restricted, due to scalability issues, to the analysis of a small number of iterations, for instance via potential functions. By contrast, while the tightness of the proposed framework depends on solving exponentially large SDPs, as in the all-scenarios approach, it allows to trade off tightness for scalability, by solving an SDP relaxation of smaller size. Hence, in some cases of sublinear convergence, our framework, by solving small-sized relaxations on a large number of iterations, improves on the all-scenarios approach, given the conservatism introduced by guessing the potential function's structure. In addition, while also untractable for a large number of iterations, the tight formulation of our framework is still computationally cheaper than the all-scenarios approach, as it grows with 2^N (as compared to n^N , with $n \geq 2$ the cardinality of the support), and possesses a strongly symmetric structure.

8.2 Preliminaries: computer-aided analysis of inexact methods.

The PEP framework introduced in Chapter 3 for the analysis of deterministic methods can be straightforwardly extended to the analysis of inexact methods, i.e., ε and σ^2 as introduced in (8.4) are deterministic perturbations rather than random variables. Indeed, in this context, one seeks worst-case performance guarantees that hold not only for all functions in a given function class, and all starting points satisfying a given initial condition, but also for the worst possible perturbation satisfying given assumptions. To obtain a tight guarantee, it thus suffices to let, in the PEP formulation, these perturbations as variables, in addition to the set $S = \{(x_i, f_i, g_i)\}_{i=0,1,\dots,N,\star}$ that encodes the iterates of the method analyzed and evaluations of the function and (sub)gradient at these iterates.

As in the deterministic setting, the resulting PEP formulations will be

SDPs, provided that the algorithm, the function class, the performance criterion, the initial conditions, and the assumptions on the perturbations, are Gram-representable, in a slightly more general way than introduced in Definition 3.1. That is, these quantities should be expressed as linear functions in (i) the function values and (ii) the scalar products between iterates, (sub)gradients, and additive perturbations, that include ε and σ^2 , as in (8.4).

Definition 8.1 (Gram-representable expression). Let $N \in \mathbb{N}$, and finite sets $S = \{(x_i, f_i, g_i)\}_{i=0,1,\dots,N,\star} \subset (\mathbb{R}^d \times \mathbb{R} \times \mathbb{R}^d)^{N+2}$, and $V = \{(\varepsilon_i, \sigma_i^2)\}_{i=0,1,\dots,N-1} \in (\mathbb{R}^d \times \mathbb{R})^N$. An algebraic expression p involving the elements of $S \cup V$ is said to be *Gram-representable* if it can be written using a finite number of functions, linear in the f_i 's and σ_j^2 , and scalar products in x_i, g_i, ε_j , $i = 0, 1, \dots, N, \star$, $j = 0, 1, \dots, N - 1$, that is,

$$\begin{aligned} &\langle g_i, g_j \rangle, \langle x_i, x_j \rangle, \langle g_i, x_j \rangle, \text{ for all } i, j = 0, 1, \dots, N, \star, \\ &\langle g_i, \varepsilon_j \rangle, \langle x_i, \varepsilon_j \rangle, \text{ for all } i = 0, 1, \dots, N, \star, j = 0, 1, \dots, N - 1, \text{ and} \\ &\langle \varepsilon_i, \varepsilon_j \rangle, \text{ for all } i, j = 0, 1, \dots, N - 1. \end{aligned}$$

Given a Gram-representable expression p , we say a pair $(F, G) \in \mathbb{R}^{2N+2} \times \mathbb{S}^{3N+4}$ satisfies p , if it is the Gram representation

$$\begin{aligned} F &:= [f_\star, f_0, \dots, f_N, \sigma_0^2, \dots, \sigma_{N-1}^2]^\top, \\ G &:= PP^\top, P := [x_\star, x_0, \dots, x_N, g_\star, g_0, \dots, g_N, \varepsilon_0, \dots, \varepsilon_{N-1}]^\top, \end{aligned} \quad (8.5)$$

of a set $S \cup V$ satisfying p . This definition of a Gram representation slightly differs from that of Chapter 3, since it includes $\varepsilon_k, \sigma_k^2, k = 0, 1, \dots, N - 1$, but also since it explicitly contains all iterates of the method, as well as x_\star and g_\star . One could, as in Chapter 3, only keep x_0 and access all other iterates as linear combinations of the entries of G . However, for the exposition of the framework to be clearer, we first keep all variables explicit. The implicit Gram-representation, used in practice, is detailed in Appendix 8.A. To access elements of G and F , given \mathbf{e}_i the i^{th} -unit vector, we define selection vectors

- $\mathbf{x}_\star = \mathbf{e}_1 \in \mathbb{R}^{3N+4}$, and $x_k = \mathbf{e}_{k+2} \in \mathbb{R}^{3N+4}, k = 0, 1, \dots, N, \mathbf{g}_\star = \mathbf{e}_{N+3} \in \mathbb{R}^{3N+4}$,
- $\mathbf{g}_\star = \mathbf{e}_{N+3} \in \mathbb{R}^{3N+4}$, and $\mathbf{g}_k = \mathbf{e}_{N+4+k} \in \mathbb{R}^{3N+4}, k = 0, 1, \dots, N$,
- $\mathbf{v}_k = \mathbf{e}_{2N+5+k} \in \mathbb{R}^{3N+4}$,

- $\mathbf{f}_\star = \mathbf{e}_1 \in \mathbb{R}^{2N+2}$, and $\mathbf{f}_k = \mathbf{e}_{k+2} \in \mathbb{R}^{2N+2}$, $k = 0, 1, \dots, N$ and $\mathbf{s}_k = \mathbf{e}_{N+3} \in \mathbb{R}^{2N+2}$, $k = 0, 1, \dots, N-1$.

These vectors allow accessing all f_i 's and σ_j^2 's ($i = 0, 1, \dots, N, \star$, $j = 0, 1, \dots, N-1$) in F , and scalar products of x_i, g_i, ε_j 's ($i = 0, 1, \dots, N, \star$, $j = 0, 1, \dots, N-1$) in G . Then, for instance, given $k = 0, 1, \dots, N-1$, we say (F, G) satisfies $\|\varepsilon_k\|^2 \leq \sigma^2$ if $\mathbf{v}_k^\top G \mathbf{v}_k \leq \sigma^2$.

For further comparison with our framework, we present an instance of PEP formulation of an inexact first-order method.

Example 8.1: Inexact GD under bounded perturbation

Consider a single iteration of the inexact gradient descent on the class $\mathcal{F}_{0,L}$ of L -smooth convex functions, under additive deterministic perturbation of bounded norm $|\sigma|$. Let the performance measure be $f(x_1) - f(x_\star)$, and the initial condition be $\|x_0 - x_\star\|^2 \leq 1$. The exact worst-case guarantee on this problem is the solution to

$$\max_{\substack{d \in \mathbb{N}, f \in \mathcal{F}_{0,L}, \\ x_0, x_\star \in \mathbb{R}^d, \\ \varepsilon_0 \in \mathbb{R}^d}} f(x_1) - f(x_\star) \text{ s.t. } \|x_0 - x_\star\|^2 \leq 1 \quad (8.6)$$

$$x_1 = x_0 - \alpha (\nabla f(x_0) + \varepsilon_0),$$

$$\nabla f(x_\star) = 0, \|\varepsilon_0\|^2 \leq \sigma^2,$$

or equivalently, to

$$\max_{\substack{F \in \mathbb{R}^3, \\ G \in \mathbb{S}_+^8}} F^\top (\mathbf{f}_1 - \mathbf{f}_\star) \quad (8.7)$$

$$\text{s.t. } (\mathbf{x}_0 - \mathbf{x}_\star)^\top G (\mathbf{x}_0 - \mathbf{x}_\star) \leq 1$$

$$F^\top (\mathbf{f}_j - \mathbf{f}_i) \geq \mathbf{g}_i^\top G (\mathbf{x}_j - \mathbf{x}_i) + \frac{1}{2L} (\mathbf{g}_i - \mathbf{g}_j)^\top G (\mathbf{g}_i - \mathbf{g}_j), \forall i, j \in \{0, 1, \star\}$$

$$(\mathbf{x}_1 - \mathbf{x}_0 + \alpha \mathbf{g}_0)^\top G (\mathbf{x}_1 - \mathbf{x}_0 + \alpha \mathbf{g}_0) = 0$$

$$\mathbf{g}_\star^\top G \mathbf{g}_\star = 0, \mathbf{v}_0^\top G \mathbf{v}_0 = \sigma^2.$$

The second constraint, in (8.7), ensures (F, G) to be $\mathcal{F}_{0,L}$ -interpolable, i.e., as stated in Proposition 2.1,

$$f_j \geq f_i + \langle \mathbf{g}_i, \mathbf{x}_j - \mathbf{x}_i \rangle + \frac{1}{2L} \|\mathbf{g}_i - \mathbf{g}_j\|^2. \quad (8.8)$$

We are now ready to present our framework for computer-aided analysis of *stochastic* optimization methods, whose simplest formulation will be closely related to PEP formulations for the analysis of inexact methods.

8.3 Framework for automated analysis of stochastic methods

This section presents a framework for computer-aided analysis of stochastic first-order methods; via the introduction of expected values as variables to PEP formulations. We first define a few problem classes on which the framework applies, before presenting a spectrum of SDP formulations whose solutions provide convergence rates on these problems.

8.3.1 Problem definition

Given $N \in \mathbb{N}$, we analyze N iterations of a method \mathcal{M} on a function class \mathcal{F} , under stochastic noise drawn from an element of a class of distributions Ω_N . Performance of \mathcal{M} is evaluated through a measure \mathcal{P} , under an initial condition C_{init} . We impose a number of assumptions on these objects.

Assumption 8.1. Let $N \in \mathbb{N}$, a method \mathcal{M} , a function class \mathcal{F} , a class of distributions $\Omega_N(A_1, A_2, B_1, B_2, C_1, C_2, D_1, D_2, E_1, \rho)$, a performance measure \mathcal{P} , and an initial condition C_{init} . We assume the following.

- **Method \mathcal{M}** [THG17a, Definition 2.11]. We consider stochastic fixed-step linear first-order methods (FSLFOM), whose k^{th} iteration on a function f is given by

$$x_{k+1} = x_0 + \sum_{j=0}^k \alpha_{k,j} (g_j + \varepsilon_j), \quad (\mathcal{M})$$

where $g_j \in \partial f(x_j)$, ε_j is a random variable, and all coefficients $\alpha_{k,j}$ are predetermined.

Denote by $S = \{(x_j, f(x_j), g_j)\}_{j=0,1,\dots,N,\star}$ the iterates $(x_j)_{j=0,1,\dots,N}$ of \mathcal{M} , and a minimizer x_\star of f , with associated function and (sub)gradient values.

- **Distribution class $\Omega_N(\mathbf{P})$.** Let $\mathbf{P} = A_1, A_2, B_1, B_2, C_1, C_2, D_1, D_2, E_1, \rho$. We consider classes of distributions $\Omega_N(\mathbf{P})$ such that any random

variable $(\varepsilon, \sigma^2) = \{(\varepsilon_i, \sigma_i^2)\}_{i=0,1,\dots,N-1}$ drawn from a distribution in $\Omega_N(\mathbf{P})$ satisfies

(i) $\mathbb{E}[\varepsilon_k | \mathcal{A}_k] = 0$, $\forall k = 0, 1, \dots, N$, where $\mathcal{A}_k = \sigma(x_0, \dots, x_k, \varepsilon_0, \dots, \varepsilon_{k-1}, \sigma_0^2, \dots, \sigma_k^2)$ is the filtration of available information at iteration k .

(ii) (ε, σ^2) satisfies (8.4) w.r.t. S , with parameters \mathbf{P} .

- **Function class \mathcal{F}** [THG17a, Definition 2.2]. We consider function classes whose interpolation conditions, see Definition 2.6, are Gram-representable on realizations of S (Definition 8.1).
- **Performance measure \mathcal{P} , and initial condition C_{init}** [THG17a, Definitions 2.3, 2.4]. We consider performance measures and initial conditions that are Gram-representable on S (Definition 8.1).

In what follows, we remove dependence of Ω_N in \mathbf{P} , to lighten notation. By an abuse of notation, we denote by $(\varepsilon, \sigma^2) \in \Omega_N$ any pair of random variables drawn from a distribution in Ω_N , whatever this specific distribution is. Observe that Assumption 8.1 allows the noise to follow different distributions from query to query, i.e., ε_k 's distribution may differ from that of ε_l , when $l \neq k$.

The proposed framework adapts to more general settings than those satisfying Assumption 8.1, e.g., there could be more randomness than simply noisy (sub)gradient evaluations, or one could consider, e.g., adaptive stepsize by modifying the filtration \mathcal{A}_k . However, Assumption 8.1 simplifies the exposition of the framework, and already encompasses a broad range of settings, including, e.g., stochastic (accelerated) gradient descent on $\mathcal{F}_{\mu, L}$, where the initial condition and performance measure may depend on $\|x_0 - x_\star\|^2$, $\|g_0\|^2$ for $g_0 \in \partial f(x_0)$, or $f(x_0) - f(x_\star)$, with analogous expressions involving x_N , g_N , and $f(x_N)$ in the performance criterion.

Central question We seek to obtain, given satisfaction of C_{init} , the tightest bound possible on the expected value of the performance measure \mathcal{P} , valid for any $f \in \mathcal{F}$ and $(\varepsilon, \sigma^2) \in \Omega_N$, that is,

$$\max_{\substack{d \in \mathbb{N}, \\ f \in \mathcal{F} \\ x_0, x_\star \in \mathbb{R}^d \\ (\varepsilon, \sigma^2) \in \Omega_N}} \mathbb{E}[\mathcal{P}] \text{ s.t. } C_{\text{init}} \text{ is satisfied,}$$

$(x_k)_{k=0,1,\dots,N}$ is generated by \mathcal{M} w.r.t. ε ,

$$0 \in \partial f(x_*) . \quad (8.9)$$

Note that x_0 and x_* (and the associated function and (sub)gradient values) are deterministic, in contrast to x_k , $k \geq 1$ (and associated function and (sub)gradient values), which are stochastic.

As an example, consider the analysis of a single iteration of (SGD) on $\mathcal{F}_{\mu,L}$, under the additive bounded noise model.

Example 8.2: SGD under additive bounded noise model

Let $N = 1$, \mathcal{M} be (SGD), $\mathcal{F} = \mathcal{F}_{\mu,L}$ for some $-\infty < \mu < L \leq +\infty$, $L \geq 0$, $\mathcal{P} = f(x_1) - f(x_*)$, $C_{\text{init}} = \|x_0 - x_*\|^2 \leq 1$, and all parameters in (8.4) be zero, except for $D_1 = \sigma^2$, for some $\sigma^2 \in \mathbb{R}$. Then, (8.9) amounts to

$$\max_{\substack{d \in \mathbb{N}, \\ f \in \mathcal{F}_{\mu,L}, \\ x_0, x_* \in \mathbb{R}^d, \\ \varepsilon_0 \text{ a r.v.}}} \mathbb{E}[f(x_1) - f(x_*)] \text{ s.t. } \|x_0 - x_*\|^2 \leq 1, \quad (8.10)$$

$$x_1 = x_0 - \alpha(\nabla f(x_0) + \varepsilon_0),$$

$$\nabla f(x_*) = 0,$$

$$\mathbb{E}[\varepsilon_0] = 0, \mathbb{E}[\varepsilon_0^2] \leq \sigma^2.$$

8.3.2 SDP relaxation of (8.9)

Exact reformulation of (8.9) As in the inexact or deterministic case, (8.9) is as such untractable, both given the infinite-dimensional variable $f \in \mathcal{F}$, the random variables $(\varepsilon, \sigma^2) \in \Omega_N$, and the expectation $\mathbb{E}[\mathcal{P}]$. We thus introduce sets of finite-dimensional variables, along with corresponding constraints that ensure maximizing over these variables is equivalent (or provides a relaxation as tight as possible) to the original problem (8.9). In particular, given $f \in \mathcal{F}$, $x_* \in \mathbb{R}^d$ a minimizer of f , $x_0 \in \mathbb{R}^d$ a starting point satisfying the initial condition C_{init} , $(\varepsilon, \sigma^2) \in \Omega_N$, and $(x_k)_{k=0,1,\dots,N}$ the set of iterates generated by \mathcal{M} accordingly, we define

$$S \cup V, \text{ where } S = \{(x_j, f(x_j), g_j)\}_{j=0,1,\dots,N,*}, g_j \in \partial f(x_j), \forall j = 0,1,\dots,N,* , \text{ and} \quad (8.11)$$

$$V = \{(\varepsilon_j, \sigma_j^2)\}_{j=0,1,\dots,N-1}.$$

Further, under Assumption 8.1, all relevant quantities in (8.9) are Gram-representable, hence it is convenient to consider the Gram representation of $S \cup V$, i.e.,

$$\begin{aligned} F &:= [f(x_\star), f(x_0), \dots, f(x_N), \sigma_0^2, \dots, \sigma_{N-1}^2]^\top, \\ G &:= PP^\top, P := [x_\star, x_0, \dots, x_N, g_\star, g_0, \dots, g_N, \varepsilon_0, \dots, \varepsilon_{N-1}]^\top. \end{aligned} \quad (8.12)$$

One cannot directly optimize over the stochastic quantity (F, G) , or over all its possible realizations, as done in [KHG23, TB19, CKPG25], since there might be an infinite number of them. Hence, following the approach of [AdKZ23b], we consider *expectations* of (F, G) as variables of interest. Specifically, let

$$(F_{\mathbb{E}}, G_{\mathbb{E}}) := (\mathbb{E}[F], \mathbb{E}[G]), \quad (8.13)$$

that thus depend on $f, x_\star, x_0, \mathcal{M}$, and (ε, σ) . Specifically, $F_{\mathbb{E}}$ contains as elements, e.g., $\mathbb{E}[f(x_i)]$, and $G_{\mathbb{E}}$ contains, e.g., $\mathbb{E}[\langle x_i, g_j \rangle]$, $\mathbb{E}[\langle x_i, \varepsilon_j \rangle]$, ...

Example 8.3: SGD under additive bounded noise model

Consider the setting of Example 8.2, i.e., the analysis of a single iteration of (SGD) on $\mathcal{F}_{\mu, L}$, under the additive bounded noise setting. Then,

$$\begin{aligned} F_{\mathbb{E}} &= [f(x_\star), f(x_0), \mathbb{E}[f(x_1)]]^\top, \\ G_{\mathbb{E}} &= \begin{pmatrix} \|x_\star\|^2 & \langle x_\star, x_0 \rangle & \dots & \mathbb{E}[\langle x_\star, g_1 \rangle] & \mathbb{E}[\langle x_\star, \varepsilon_0 \rangle] \\ \langle x_\star, x_0 \rangle & \|x_0\|^2 & \dots & \mathbb{E}[\langle x_0, g_1 \rangle] & \mathbb{E}[\langle x_0, \varepsilon_0 \rangle] \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ \mathbb{E}[\langle x_\star, g_1 \rangle] & \mathbb{E}[\langle x_0, g_1 \rangle] & \dots & \mathbb{E}[\|g_1\|^2] & \mathbb{E}[\langle g_1, \varepsilon_0 \rangle] \\ \mathbb{E}[\langle x_\star, \varepsilon_0 \rangle] & \mathbb{E}[\langle x_0, \varepsilon_0 \rangle] & \dots & \mathbb{E}[\langle g_1, \varepsilon_0 \rangle] & \mathbb{E}[\|\varepsilon_0\|^2] \end{pmatrix}. \end{aligned}$$

Our goal is to maximize the expectation of the performance criterion \mathcal{P} , which is, under Assumption 8.1, linear in realizations of (F, G) . Equivalently, by linearity of the expectation, it suffices to maximize the performance criterion over all possible $(F_{\mathbb{E}}, G_{\mathbb{E}})$, generated by any $f \in \mathcal{F}$, $(\varepsilon, \sigma) \in \Omega_N$, x_0 and x_\star . We therefore introduce the set $\mathcal{G}_{\mathbb{E}}(\mathcal{M}, \mathcal{F}, \Omega_N, \mathcal{C}_{\text{init}})$

of pairs (F, G) that correspond to some (F_E, G_E) , i.e.,

$$\mathcal{G}_E := \left\{ \begin{array}{l} (F, G) \in \mathbb{R}^{N+2} \times \mathbb{S}_+^{4N+4} : \exists (\varepsilon, \sigma) \in \Omega_N, f \in \mathcal{F}, \\ \quad x_* \in \mathbb{R}^d \text{ a minimizer of } f, \\ \quad x_0 \in \mathbb{R}^d \text{ satisfying } C_{\text{init}}, \text{ and} \\ \quad (x_k)_{k=0,1,\dots,N} \text{ the iterates of } \mathcal{M} \\ \text{s.t. } (F, G) = (F_E, G_E) \end{array} \right\}, \quad (8.14)$$

where the dependence of \mathcal{G}_E on $\mathcal{M}, \mathcal{F}, \Omega_N$ and C_{init} is removed and supposed clear from the context. Then, (8.9) can be reformulated as

$$\max_{(F,G) \in \mathcal{G}_E} \text{Tr}(A_{\text{obj}}G) + F^\top b_{\text{opt}}, \quad (P_{\text{exact}})$$

where the linear terms $\text{Tr}(A_{\text{obj}}G)$ and $F^\top b_{\text{opt}}$ encode the performance measure \mathcal{P} .

Example 8.4: SGD under additive bounded noise model

Consider the setting of Example 8.2, i.e., the analysis of a single iteration of (SGD) on $\mathcal{F}_{\mu,L}$, under the additive bounded noise setting. Then, (P_{exact}) becomes

$$\max_{(F,G) \in \mathcal{G}_E((\text{SGD}), \mathcal{F}_{\mu,L}, \Omega_N(D_1 = \sigma^2), \|x_0 - x_*\|^2 \leq 1)} F^\top (\mathbf{f}_1 - \mathbf{f}_*)$$

where $\mathbf{f}_1 = \mathbf{e}_3 \in \mathbb{R}^3$ and $\mathbf{f}_* = \mathbf{e}_1 \in \mathbb{R}^1$ are selection vectors accessing $f(x_1)$, $f(x_*)$ in F .

To solve (P_{exact}) , our goal is to represent \mathcal{G}_E , using a finite number of algebraic conditions. As in [AdKZ23b], we first introduce conditions that are *necessary*, but not sufficient, for inclusion in \mathcal{G}_E .

As a reminder, given a Gram representable condition p (Definition 8.1), we say (F, G) satisfies p if it is the Gram representation (8.5) of some $S = \{(x_i, f_i, g_i)\}_{i=0,1,\dots,N,*}$ and $V = \{(\varepsilon_i, \sigma_i^2)\}_{i=0,1,\dots,N-1}$ satisfying p , see Appendix 8.A for details. The necessary conditions describing \mathcal{G}_E , given $\mathcal{M}, \mathcal{F}, \Omega_N$ and C_{init} satisfying Assumption 8.1, follow from

- Gram-representable constraints that necessarily hold on any realization, imposed in expectation—namely, any $(F, G) \in \mathcal{G}_E$ should be \mathcal{F} -

interpolable, satisfy the initial condition C_{init} , the method definition (\mathcal{M}) , and the optimality condition $\|g_\star\|^2 = 0$; and

- Gram-representable constraints that naturally involve expectations. Specifically, these include (i) the noise variance (8.4), that is, $(F, G) \in \mathcal{G}_E$ should satisfy, given constants $A_1, A_2, B_1, B_2, C_1, C_2, D_1, D_2, \rho \in \mathbb{R}$, and $\forall k = 0, 1, \dots, N-1$,

$$\begin{aligned} \|\varepsilon_k\|^2 &\leq A_1(f_k - f_\star) + B_1\|g_k\|^2 + C_1\|x_k - x_\star\|^2 + D_1 + E_1\sigma_k^2, \\ \sigma_{k+1}^2 &\leq A_2(f_k - f_\star) + B_2\|g_k\|^2 + C_2\|x_k - x_\star\|^2 + D_2 + (1 - \rho)\sigma_k^2, \end{aligned}$$

(C_{variance})

as well as (ii) uncorrelation conditions involving ε_k and ε_l ($l \neq k$), and ε_k and x_i, g_i ($i \leq k$). Specifically, $(F, G) \in \mathcal{G}_E$ should satisfy

$$\langle \varepsilon_k, \varepsilon_l \rangle = \langle \varepsilon_k, x_i \rangle = \langle \varepsilon_k, g_i \rangle = 0, \quad \forall k = 0, 1, \dots, N-1, i \leq k, l \neq k. \quad (\text{C}_\perp)$$

For instance, condition $\langle \varepsilon_k, \varepsilon_l \rangle = 0$ writes $\mathbf{v}_k^\top \mathbf{G} \mathbf{v}_l = 0$. Hence, if $G = G_E$, it becomes $\mathbb{E}[\langle \varepsilon_k, \varepsilon_l \rangle] = 0$, which holds by definition of Ω_N . We now formalize the proposed conditions.

Theorem 8.1 (Properties of expected variables)

Given $N \in \mathbb{N}$, consider a method \mathcal{M} , a function class \mathcal{F} , a class of distributions Ω_N , and an initial condition C_{init} , satisfying Assumption 8.1. In addition, let \mathcal{G}_E be defined in (8.14). Then, any $(F, G) \in \mathcal{G}_E$ is \mathcal{F} -interpolable, and satisfies $C_{\text{init}}, C_{\text{variance}}, C_\perp, \|g_\star\|^2 = 0$, and (\mathcal{M}) .

Proof. Let $(F_E, G_E) \in \mathcal{G}_E$. Denote by (\tilde{F}, \tilde{G}) the stochastic pair for which $F = \mathbb{E}[\tilde{F}], G = \mathbb{E}[\tilde{G}]$.

By definition of \mathcal{G}_E , given in (8.14), there exist some $(\varepsilon, \sigma^2) \in \Omega_N, f \in \mathcal{F}, x_\star \in \mathbb{R}^d$ minimizer of $f, x_0 \in \mathbb{R}^d$ satisfying C_{init} , and $(x_k)_{k=0,1,\dots,N}$ the associated iterates of \mathcal{M} , such that (\tilde{F}, \tilde{G}) is the Gram representation (8.12) of $S = \{(x_j, f(x_j), g_j)\}_{j=0,1,\dots,N,\star}$, where $g_j \in \partial f(x_j)$, and $V = \{(\varepsilon_j, \sigma_j^2)\}_{j=0,1,\dots,N-1}$. Hence, any realization of (\tilde{F}, \tilde{G}) is \mathcal{F} -interpolable, and satisfies $C_{\text{init}}, \|g_\star\|^2 = 0$ and (\mathcal{M}) . By linearity of these conditions in the realizations of (\tilde{F}, \tilde{G}) , and linearity of an expectation, it therefore holds that (F_E, G_E) is \mathcal{F} -interpolable and satisfies $C_{\text{init}}, \|g_\star\|^2 = 0$ and (\mathcal{M}) .

Further, satisfaction of C_{variance} by (F_E, G_E) follows straightforwardly from satisfaction of (8.4) by $(\varepsilon, \sigma^2) \in \Omega_N$. In addition, $\langle \varepsilon_k, \varepsilon_l \rangle = 0$, $k \neq l$, and $\langle \varepsilon_k, x_i \rangle = \langle \varepsilon_k, g_i \rangle = 0$, $i \leq k$, follows from $\mathbb{E}[\varepsilon_k | \mathcal{A}_k] = 0$. Indeed, then, e.g., by the law of total expectation, $\mathbb{E}[\langle \varepsilon_k, x_i \rangle] = \mathbb{E}[\mathbb{E}[\langle \varepsilon_k, x_i \rangle | \mathcal{A}_k]] = \mathbb{E}[\langle \mathbb{E}[\varepsilon_k | \mathcal{A}_k], x_i \rangle] = 0$, and the same holds for the other variables in \mathcal{A}_k . \square

The necessary conditions introduced in Theorem 8.1 differ from the ones proposed in [AdKZ23b], under the coordinate descent setting, in the condition C_{\perp} . Indeed, in [AdKZ23b], the only uncorrelation conditions imposed are $\langle \varepsilon_k, x_k \rangle = \langle \varepsilon_k, g_k \rangle = 0$, $k = 0, 1, \dots, N-1$. Figure 5.5 illustrates, on a numerical example, the difference between the two formulations.

Simplest SDP relaxation of (8.9) Relying on Theorem 8.1, we introduce a set \mathcal{G}_c satisfying $\mathcal{G}_E \subseteq \mathcal{G}_c$.

$$\mathcal{G}_c := \left\{ (F, G) \in \mathbb{R}^{2N+2} \times \mathbb{S}_+^{3N+4} \mid \begin{array}{l} (F, G) \text{ is } \mathcal{F}\text{-interpolable, and} \\ \text{satisfies } C_{\text{init}}, (\mathcal{M}), \|g_{\star}\|^2 = 0, \\ C_{\perp} \text{ and } C_{\text{variance}} \end{array} \right\}. \quad (8.15)$$

Under Assumption 8.1, this set is convex, since characterized by a set of linear conditions on (F, G) . In addition, it allows for a SDP relaxation of (8.9), given by

$$\max_{(F, G) \in \mathcal{G}_c} \text{Tr}(A_{\text{obj}}G) + F^{\top} b_{\text{opt}}, \quad (P_{\text{relax},1})$$

where the linear terms $\text{Tr}(A_{\text{obj}}G)$ and $F^{\top} b_{\text{opt}}$ encode the performance measure \mathcal{P} . Indeed, by Theorem 8.1, $\mathcal{G}_E \subseteq \mathcal{G}_c$, hence the bound resulting from (P_{exact}) is smaller than the one resulting from $(P_{\text{relax},1})$.

Example 8.5: SGD under additive bounded noise model

Consider the setting of Example 8.2, i.e., a single iteration of (SGD) on $\mathcal{F}_{\mu, L}$, under additive bounded noise. Then, $(P_{\text{relax},1})$ can be expressed as

$$\max_{\substack{F \in \mathbb{R}^3, \\ G \in \mathbb{S}_+^6}} F^{\top} (\mathbf{f}_1 - \mathbf{f}_{\star}) \quad (8.16)$$

$$\begin{aligned}
 & \text{s.t. } (\mathbf{x}_0 - \mathbf{x}_\star)^\top G(\mathbf{x}_0 - \mathbf{x}_\star) \leq 1 \\
 & F^\top (\mathbf{f}_j - \mathbf{f}_i) \geq \mathbf{g}_i^\top G(\mathbf{x}_j - \mathbf{x}_i) + \frac{1}{2L} (\mathbf{g}_i - \mathbf{g}_j)^\top G(\mathbf{g}_i - \mathbf{g}_j), \forall i, j \in \{0, 1, \star\} \\
 & (\mathbf{x}_1 - \mathbf{x}_0 + \alpha \mathbf{g}_0)^\top G(\mathbf{x}_1 - \mathbf{x}_0 + \alpha \mathbf{g}_0) = 0 \\
 & \mathbf{g}_\star^\top G \mathbf{g}_\star = 0, \mathbf{v}_0^\top G \mathbf{v}_0 = \sigma^2. \\
 & \mathbf{v}_0^\top G \mathbf{g}_0 = \mathbf{v}_0^\top G \mathbf{x}_0 = 0.
 \end{aligned} \tag{8.17}$$

The formulation in (8.16) is equivalent to (8.7), except for the independence conditions. The key difference lies in the interpretation of the variables: while (8.7) operates on the (deterministic) Gram representation of $S \cup V$, the variables (F, G) in (8.16) represent the *expectation* of this Gram representation in the stochastic setting. As a result, even when relaxing the exact stochastic formulation (8.9), working with expectations directly allows for potentially tighter worst-case bounds than those derived from deterministic analyses.

Remark 8.1. Although the variables $(\varepsilon_k, \sigma_k^2)$ represent stochastic noise, \mathcal{G}_c interprets them as deterministic entities satisfying (8.4), and orthogonal to all prior variables x_i, g_i , where $i \leq k$, and $\varepsilon_i, i < k$.

Elements $(G, F) \in \mathcal{G}_c$ are intended to represent expectations of the Gram representation of the set $S \cup V$ as defined in (8.11), but \mathcal{G}_c only encodes necessary conditions for this to hold. In particular, \mathcal{G}_c may include pairs (G, F) that do not correspond to any valid expected Gram representation. This situation is analogous to the case where only necessary conditions for \mathcal{F} -interpolability are imposed, leading to sets that may contain elements not consistent with any function in \mathcal{F} .

Due to its scalability, we will mostly use the relaxation $(P_{\text{relax},1})$ to analyze stochastic methods, in Section 8.4. However, to provide a complete framework, that allows obtaining tight performance guarantees, we first present a sequence of refinements of the conditions describing \mathcal{G}_E , eventually leading to sufficient conditions, that yield increasingly tighter relaxations of (8.9). We first describe some of the sources of relaxation in $(P_{\text{relax},1})$, arising from the difference between \mathcal{G}_E and \mathcal{G}_c .

Sources of relaxation in $(P_{\text{relax},1})$ Suppose that $(G, F) \in \mathcal{G}_c$ indeed corresponds to an expectation with respect to some stochastic noise $(\varepsilon, \sigma^2) \in \Omega_N$. Still, it may not belong to \mathcal{G}_E , for several reasons.

Compensation across realizations. In \mathcal{G}_E , the constraints C_{init} , $\|g_\star\| = 0$, (\mathcal{M}) , and \mathcal{F} -interpolability are satisfied by the union of all realizations of (F, G) , generated by a noise realization $\varepsilon^{(i)}$. By contrast, in \mathcal{G}_c , these conditions are only required to hold in expectation, allowing for compensation across different realizations. As a result, some realizations may individually violate one of these constraints, while their average satisfies them.

Realizations interpolated by different functions. Even if each realization of (F, G) is \mathcal{F} -interpolable, they may be consistent with different functions in \mathcal{F} , rather than a single common one. Belonging to \mathcal{G}_E requires that a single function $f \in \mathcal{F}$ simultaneously interpolates all realizations.

Example 8.1. As a toy example, consider a slightly modified setting: let

- $x_0 = \varepsilon_0$ a random variable with 2 equiprobable realizations, -1 and 1 ,
- \mathcal{F} the class of convex functions f with bounded subgradients, that is, $\|g_x\| \leq 1$, $\forall g_x \in \partial f(x)$, $\forall x \in \mathbb{R}^d$, and
- a method $x_1 = x_0 - 2g_0$, where $g_0 \in \partial f(x_0)$.

Then,

$$\begin{aligned} \begin{pmatrix} x_0 \\ x_1 \end{pmatrix} &= \begin{pmatrix} -1 \\ 1 \end{pmatrix}, \quad \begin{pmatrix} g_0 \\ g_1 \end{pmatrix} = \begin{pmatrix} -1 \\ 0 \end{pmatrix}, \quad \begin{pmatrix} f_0 \\ f_1 \end{pmatrix} = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \\ \begin{pmatrix} \tilde{x}_0 \\ \tilde{x}_1 \end{pmatrix} &= \begin{pmatrix} 1 \\ -1 \end{pmatrix}, \quad \begin{pmatrix} \tilde{g}_0 \\ \tilde{g}_1 \end{pmatrix} = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \quad \begin{pmatrix} \tilde{f}_0 \\ \tilde{f}_1 \end{pmatrix} = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \end{aligned}$$

satisfy

$$\begin{aligned} x_1 &= x_0 - 2g_0, \quad \tilde{x}_1 = \tilde{x}_0 - 2\tilde{g}_0 \\ \|g_0\|, \|g_1\|, \|\tilde{g}_0\|, \|\tilde{g}_1\| &\leq 1 \\ f_0 - f_1 &\geq g_1(x_0 - x_1), \quad f_1 - f_0 \geq g_0(x_1 - x_0) \\ \tilde{f}_0 - \tilde{f}_1 &\geq \tilde{g}_1(\tilde{x}_0 - \tilde{x}_1), \quad \tilde{f}_1 - \tilde{f}_0 \geq \tilde{g}_0(\tilde{x}_1 - \tilde{x}_0). \end{aligned}$$

That is, the pairs (F, G) and (\tilde{F}, \tilde{G}) generated by both realizations of x_0 are consistent with the method and \mathcal{F} -interpolable. By linearity, their average $\frac{1}{2}(F, G) + \frac{1}{2}(\tilde{F}, \tilde{G})$ belongs to \mathcal{G}_c . However, it does not belong to

\mathcal{G}_E . Indeed, it is the solution to $\max_{(F,G) \in \mathcal{G}_c} f_0 - f_1$, which differs from $\max_{(F,G) \in \mathcal{G}_E} f_0 - f_1$, as computed exactly via the formulation introduced in Theorem 8.3. This furnishes an instance of \mathcal{G}_c suffering from the second source of conservatism.

8.3.3 Refinement of $(P_{\text{relax},1})$: a multi-realizations approach

Attempting to close the gap between \mathcal{G}_c and \mathcal{G}_E , we refine the description of \mathcal{G}_E introduced in Theorem 8.1.

Overview: from copies to scenarios The proposed refinement relies on the introduction of several copies (F^i, G^i) of any $(F, G) \in \mathcal{G}_E$, such that optimizing over a single copy or the average of such copies is equivalent. One can then relax the optimization over these copies belonging to \mathcal{G}_E onto an optimization over several pairs belonging to (G, C) , which thus provides an upper bound on (8.9). We then show that for wisely chosen copies, the associated instances in \mathcal{G}_c can be interpreted as scenarios generated by a noise with specific equiprobable two-realizations distribution. The associated guarantee is thus, also, a lower bound on (8.9). This allows (i) obtaining an exact SDP formulation of (8.9), and (ii) identifying the worst-distribution within a noise model. The exact SDP formulation is thus, in its implementation, exactly an all-scenarios approach, with 2^N scenarios of a specific equiprobable 2-realizations distribution. But, it differs in interpretation, since the resulting guarantees hold not only for this specific distribution, but for all distributions with same variance and expectation.

Refinement of \mathcal{G}_c Informally, let $(F, G) \in \mathcal{G}_E$, generated by $(\varepsilon, \sigma^2) \in \Omega_N$, $f \in \mathcal{F}$, $x_* \in \mathbb{R}^d$, and $x_0 \in \mathbb{R}^d$. Then, one can consider R copies $(\varepsilon^i, \sigma^{2i})$ of (ε, σ^2) , identically distributed to (ε, σ^2) , but that can be correlated. From these copies, and using the same f, x_0, x_* , one can generate copies $(F^i, G^i) \in \mathcal{G}_E$ of (F, G) . These copies are by definition identical, hence solving (P_{exact}) is equivalent to maximizing the expected performance measure over their average, i.e.,

$$\max_{\substack{(F^i, G^i) \in \mathcal{G}_E \\ \text{copies generated by identically distributed } (\varepsilon^i, \sigma^{2i}) \in \Omega_N \\ \text{and same } f \in \mathcal{F}, x_0, x_* \in \mathbb{R}^d}} \frac{1}{R} \sum_i \text{Tr} (A_{\text{obj}} G^i) + (F^i)^\top b_{\text{opt}}. \tag{8.18}$$

Formally, equivalence of (8.18) and (P_{exact}) follows from Proposition 8.1.

Proposition 8.1. *Consider a random variable v , a function $P(v)$, and $R \in \mathbb{N}$. Then,*

$$\mathbb{E}_v[P(v)] = \frac{1}{R} \mathbb{E}_{v^1, \dots, v^R} \left[\sum_{i=1}^R P(v^i) \right], \quad (8.19)$$

where v^i , $i = 1, 2, \dots, R$ follow the same distribution as v , and may be correlated.

Proof. It holds that $\mathbb{E}_v[P(v)] = \frac{1}{R} (\sum_{i=1}^R \mathbb{E}_{v^i}[P(v^i)]) = \frac{1}{R} \mathbb{E}_{v^1, \dots, v^R} [\sum_{i=1}^R P(v^i)]$, by linearity. \square

Similarly to the relaxation of (P_{exact}) to $(P_{\text{relax},1})$, maximizing over \mathcal{G}_c rather than \mathcal{G}_E , (8.18) can be relaxed by maximizing the expected performance measure over the average of R elements in \mathcal{G}_c , meant to correspond to R identical copies in \mathcal{G}_E . These elements should necessarily (i) share the same deterministic quantities $(x_i, g_i, f_i)_{i=0, \star}$, and (ii) be interpolated by a common function in f , i.e., they should satisfy interpolation constraints in between elements.

In addition, the advantage of maximizing the expected performance measure over the average of identical copies, rather than over a single copy, lies in the fact that the copies may be correlated. Regardless of the correlation, these two maximization problems are equivalent when the variables belong to \mathcal{G}_E . However, when they only belong to \mathcal{G}_c , wisely choosing this correlation allows strengthening $(P_{\text{relax},1})$ by introducing covariance constraints linking the elements in \mathcal{G}_c . In particular, if these elements are chosen to be uncorrelated, no improvement is obtained compared to the initial relaxation $(P_{\text{relax},1})$. Else, the induced relationships reduce the sources of relaxation described in Section 8.3.2. Given $N, R \in \mathbb{N}$, we encode this covariance as a set of diagonal $N \times N$ matrices

$$\Sigma_{N,R} = (\Sigma_{N,R}^{ij})_{i,j \in \{1,2,\dots,R\}, j > i}, \text{ where } \Sigma_{N,R}^{ij} \in \mathcal{D}^{N \times N} \text{ satisfies } \|\Sigma_{N,R}^{ij}\| \leq 1. \quad (8.20)$$

Then, we relax (8.18) as a maximization over an average of R elements in \mathcal{G}_c , interpolated by a single function in \mathcal{F} , sharing identical deterministic quantities $(x_i, g_i, f_i)_{i=0, \star}$, and satisfying correlation conditions defined by

$\Sigma_{N,R}$, that is,

$$\max_{\substack{\{(F^i, G^i)\}_{i=1, \dots, R} \\ G^{ij} \in \mathbb{S}^{3N+4}, j=2, \dots, R}} \frac{1}{R} \sum_{i=1}^R \text{Tr} (A_{\text{obj}} G^i) + (F^i)^\top b_{\text{opt}}, \quad (P_{\text{relax, multi}})$$

$$\text{s.t. } (F^i, G^i) \in \mathcal{G}_c,$$

$$(\mathbf{F}, \mathbf{G}) \text{ is } \mathcal{F}\text{-interpolable, } \mathbf{G} \succcurlyeq 0, \quad (8.21)$$

$$\text{where } (\mathbf{F}, \mathbf{G}) = \left(\begin{bmatrix} F^1 \\ F^2 \\ \vdots \\ F^R \end{bmatrix}, \begin{pmatrix} G^1 & G^{12} & \dots & G^{1R} \\ G^{12} & G^2 & \dots & G^{2R} \\ \vdots & \vdots & \ddots & \vdots \\ G^{1R} & G^{2R} & \dots & G^R \end{pmatrix} \right)$$

$$\forall i, j \in \{1, \dots, R\}, \|x_0^i\|^2 = \|x_0^j\|^2, \|g_0^i\|^2 = \|g_0^j\|^2, f_0^i = f_0^j,$$

$$\|x_\star^i\|^2 = \|x_\star^j\|^2, f_\star^i = f_\star^j,$$

$$\forall j > i, \forall k = 0, 1, \dots, N-1, \langle \varepsilon_k^i, \varepsilon_k^j \rangle = (\Sigma_{N,R}^{ij})_{k+1, k+1} \min_l \|\varepsilon_k^l\|^2,$$

where $(\Sigma_{N,R}^{ij})_{k,l}$ is the (k,l) th entry of $\Sigma_{N,R}^{ij}$. The minimization over all $\|\varepsilon_k^l\|$, in the last constraint in $(P_{\text{relax, multi}})$ ensures that $\langle \varepsilon_k^i, \varepsilon_k^j \rangle \leq \|\varepsilon_k^i\| \|\varepsilon_k^j\|$, and can be implemented via a set of R Gram-representable inequalities. When $\|\varepsilon_k^i\| = \|\varepsilon_k^j\|$ for all $i, j = 1, 2, \dots, R$, e.g., under the additive bounded noise model, this minimization can be removed.

Note that, again, all constraints in $(P_{\text{relax, multi}})$ are Gram-representable. Hence, for the sake of clarity, we express the constraints in terms of copies $S^i \cup V^i = \{(x_k^i, f_k^i, g_k^i)\}_{k=0,1, \dots, N, \star} \cup \{(\varepsilon_k^i, \sigma_k^{2i})\}_{k=0,1, \dots, N-1}$ of which (F^i, G^i) are the Gram representation. The transition from one to the other is however straightforward, e.g., $\langle \varepsilon_k^i, \varepsilon_k^j \rangle = \mathbf{v}_k^\top G^{ij} \mathbf{v}_k$, where \mathbf{v}_k is the selection vector accessing the element ε_k in G , see Appendix 8.A for details.

Example 8.6: SGD under additive bounded noise model

Let $R = 2$, and consider Example 8.2, i.e., a single iteration of (SGD) on the class $\mathcal{F} = \mathcal{F}_{0, \text{inf}}$ of convex functions, under additive bounded

noise. Let $\Sigma_{1,2}^{12} = -1$. Then, $(P_{\text{relax,multi}})$ becomes

$$\begin{aligned} & \max_{\substack{\{(F^i, G^i)\}_{i=1,2} \in \mathbb{R}^3 \times S_+^6 \\ G^{12} \in S^6}} \frac{1}{2} \sum_{i=1}^2 (F^i)^\top (\mathbf{f}_1 - \mathbf{f}_\star) \quad (8.22) \\ & \text{s.t. } \begin{pmatrix} G^1 & G^{12} \\ G^{12} & G^2 \end{pmatrix} \succcurlyeq 0, \end{aligned}$$

\mathcal{F} -interpolability: $f_i^k \geq f_j^l + \langle g_j^l, x_i^k \rangle - \langle g_j^l, x_j^l \rangle$, $k, l = 1, 2$, $i, j \in \{0, 1, \star\}$,

C_{init} : $\|x_0^1 - x_\star^1\| \leq 1$,

$C_{\perp}, C_{\text{variance}}$: $\langle \varepsilon_0^i, x_0^i \rangle = \langle \varepsilon_0^i, g_0^i \rangle = 0$, $\|\varepsilon_0^i\|^2 = \sigma^2$, $i = 1, 2$,

Same x_0 and x_\star : $\|x_0^1\|^2 = \|x_0^2\|^2$, $\|g_0^1\|^2 = \|g_0^2\|^2$, $f_0^1 = f_0^2$,

$\|x_\star^1\|^2 = \|x_\star^2\|^2$, $f_\star^1 = f_\star^2$,

Correlation: $\langle \varepsilon_0^1, \varepsilon_0^2 \rangle = -\sigma^2$,

where the interpolation conditions follow from (8.8), imposed (i) on pairs x_i, x_j from a single element in \mathcal{G}_c ($k = l = 1$), or (ii) on pairs x_i, x_j from two elements in \mathcal{G}_c ($k \neq l$). The difference with (8.16) lies in imposing a single function to interpolate two elements in \mathcal{G}_c , linked via the correlation $\langle \varepsilon_0^1, \varepsilon_0^2 \rangle = -\sigma^2$, that prevents these elements to be chosen identical.

We now show that $(P_{\text{relax,multi}})$ improves upon $(P_{\text{relax},1})$, while remaining a relaxation of (P_{exact}) .

Theorem 8.2 (Spectrum of SDP formulations)

Given $N, R \in \mathbb{N}$, consider a method \mathcal{M} , a function class \mathcal{F} , a class of distributions Ω_N , and an initial condition C_{init} , satisfying Assumption 8.1. In addition, consider a covariance matrix $\Sigma_{N,R}$ satisfying (8.20). Let f_{obj}^\star be the solution to $(P_{\text{relax,multi}})$. Then,

$$\max_{(F,G) \in \mathcal{G}_E} \text{Tr}(A_{\text{obj}}G) + F^\top b_{\text{opt}} \leq f_{\text{obj}}^\star \leq \max_{(F,G) \in \mathcal{G}_c} \text{Tr}(A_{\text{obj}}G) + F^\top b_{\text{opt}},$$

where \mathcal{G}_E is defined in (8.14), and \mathcal{G}_c is defined in (8.15).

Proof. We prove the first inequality. Given $(F, G) \in \mathcal{G}_E$, let $(\varepsilon, \sigma^2) \in \Omega_N$,

$f \in \mathcal{F}$, $x_0 \in \mathbb{R}^d$ and $x_* \in \mathbb{R}^d$ be its generating random variables, function, starting point and minimizer. For any random variable (v^i, σ^{2i}) with same distribution as ε , consider $(F^i, G^i) \in \mathcal{G}_E$, the copy of (F, G) generated by (v^i, σ^{2i}) and f , with respect to the same deterministic variables. By Proposition 8.1, applied to the objective value $\text{Tr}(A_{\text{obj}}G) + F^\top b_{\text{opt}}$ of (P_{exact}) , which is the expectation of some function of ε , it holds

$$\text{Tr}(A_{\text{obj}}G) + F^\top b_{\text{opt}} = \frac{1}{R} \sum_{i=1}^R \text{Tr}(A_{\text{obj}}G^i) + (F^i)^\top b_{\text{opt}}.$$

Hence, (P_{exact}) is equivalent to (8.18). All feasible variables $(F^i, G^i) \in \mathcal{G}_E$ of (8.18) (i) belong to \mathcal{G}_c , by Theorem 8.1, (ii) are interpolated by a single function $f \in \mathcal{F}$, hence satisfy (\mathbf{F}, \mathbf{G}) as defined in $(P_{\text{relax, multi}})$ is \mathcal{F} -interpolable, and (iii) share the same deterministic variables $\{(x_i, f_i, g_i)\}_{i=0, \star}$. In addition, since v^i , $i = 1, 2, \dots, R$, are not independent, the covariance between v_k^i and v_k^j , for $k = 0, 1, \dots, N - 1$, can be chosen to be any value C_k satisfying

$$C_k \leq \sqrt{\mathbb{E}_{v^i}[(v_k^i)^2] \mathbb{E}_{v^j}[(v_k^j)^2]},$$

since by Cauchy-Schwarz inequality, this quantity bounds $E[\langle v_k^i, v_k^j \rangle]$, see, e.g., [BT08, Section 4.2]. By definition of Σ^{ij} , the covariance introduced in $(P_{\text{relax, multi}})$ satisfies this assumption. Finally, σ^{2i} , $i = 1, 2, \dots, R$, can be chosen to be uncorrelated. Hence, any feasible variable to (P_{exact}) is also a feasible variable to $(P_{\text{relax, multi}})$.

We now prove the second inequality. By linearity of the objective, $(P_{\text{relax, multi}})$ is equivalent to

$$\begin{aligned} & \max_{(F, G)} \text{Tr}(A_{\text{obj}}G) + (F)^\top b_{\text{opt}}, \\ & \text{s.t. } \exists (F^i, G^i) \in \mathcal{G}_c : (F, G) = \frac{1}{R} \sum_{i=1}^R (F^i, G^i), \end{aligned}$$

$$\text{and additional constraints on } (F^i, G^i). \quad (8.23)$$

Hence, the solution to (8.23) is smaller or equal to that of

$$\max_{(F,G)} \text{Tr} (A_{\text{obj}}G) + (F)^\top b_{\text{opt}}, \text{ s.t. } \exists (F^i, G^i) \in \mathcal{G}_c : (F, G) = \frac{1}{R} \sum_{i=1}^R (F^i, G^i). \quad (8.24)$$

By convexity of \mathcal{G}_c , any feasible variable to (8.24) belongs to \mathcal{G}_c , hence is also a feasible variable of $(P_{\text{relax},1})$, which proves the second inequality. \square

8.3.4 Using symmetry to decrease the complexity of $(P_{\text{relax},\text{multi}})$

This section proposes a class of covariance matrices $\Sigma_{N,R}$ for which $(P_{\text{relax},\text{multi}})$ admits a more compact formulation. Our goal is to introduce symmetry in the covariance constraints and render instances (F^i, G^i) in $(P_{\text{relax},\text{multi}})$ interchangeable, in the sense that swapping them yields an identical result. Indeed, then, we will be able to impose (F^i, G^i) to be identical, for all $i = 1, 2, \dots, R$, hence to (i) reduce the complexity of $(P_{\text{relax},\text{multi}})$, and (ii) remove the minimization over all $\|\varepsilon_k^l\|^2$ in the correlation constraint in $(P_{\text{relax},\text{multi}})$, since all of these quantities become equivalent. We therefore define the notion of a *super-symmetric matrix*.

Definition 8.2 (Order p super-symmetric matrix). Let $k \in \mathbb{N}$, with $p \in \mathbb{N}$, $p \geq 1$. Consider a matrix $M \in (\mathbb{S}^k)^{2^p \times 2^p}$ whose entries are themselves symmetric matrices in \mathbb{S}^k . We define order- p super-symmetry of M recursively.

- If $p = 1$, then, M is order-1 super-symmetric if, for some $A, B \in \mathbb{S}^k$,

$$M = \begin{pmatrix} A & B \\ B & A \end{pmatrix}.$$

- If $p > 1$, partition M into four blocks in $(\mathbb{S}^k)^{2^{p-1} \times 2^{p-1}}$:

$$M = \begin{pmatrix} \bar{M}_{1,1} & \bar{M}_{1,2} \\ \bar{M}_{2,1} & \bar{M}_{2,2} \end{pmatrix}.$$

Then, M is order- p super-symmetric if (i) the 2×2 block matrix $\bar{M} = (\bar{M}_{i,j})$ is order-1 super-symmetric, and (ii) each block $\bar{M}_{i,j}$ is order- $(p - 1)$ super-symmetric (if $p - 1 \geq 1$).

Example 8.2. Let $k \in \mathbb{N}$, $p = 2$. Then, $M \in (\mathbb{S}^k)^{2^p \times 2^p}$ is super-symmetric if

$$M = \begin{pmatrix} A & B & C & D \\ B & A & D & C \\ C & D & A & B \\ D & C & B & A \end{pmatrix},$$

where $A, B, C, D \in \mathbb{S}^k$.

Before proceeding to defining classes of covariance matrices, we state a few properties of super-symmetric matrices.

Proposition 8.2. *Given $k \in \mathbb{N}$ and $p \in \mathbb{N}$, $p \geq 1$, consider a super-symmetric matrix $M \in (\mathbb{S}^k)^{2^p \times 2^p}$. Then,*

- The $(2^p)^2$ entries of M all take their values among at most 2^p elements, and
- Given a permutation $\pi_p(P)$ of the vector $P = [1, 2, \dots, 2^p]$, denote by $M_{\pi_p(P)}$ the matrix obtained from M by swapping the indices of its entries according to $\pi_p(P)$.

Then, there exists 2^p permutations $\pi_p^i(P)$ such that $M_{\pi_p^i(P)} = M$, obtained by recursion:

$$\begin{aligned} -\pi_1^1([1, 2]) &= \{1, 2\}, \text{ and } \pi_1^2([1, 2]) = [2, 1] \\ -\pi_p^i([1, 2, \dots, 2^p]) &= \pi_1^i([\pi_{p-1}^i[1, \dots, 2^{p-1}], \pi_{p-1}^i[2^{p-1} + 1, \dots, 2^p]]), \\ \pi_p^{p-1+i}([1, 2, \dots, 2^p]) &= \pi_1^i([\pi_{p-1}^i[1, \dots, 2^{p-1}], \pi_{p-1}^i[2^{p-1} + 1, \dots, 2^p]]). \end{aligned}$$

Proof. We proceed by induction.

- By definition of super-symmetry, if $p = 1$, M contains at most two elements, that is, $M_{1,1}$ and $M_{1,2}$. In addition, $M = M_{[1,2]} = M_{[2,1]}$.
- Suppose the assertions are valid for p , and consider the case $p + 1$. Partition M into four blocks in $(\mathbb{S}^k)^{2^p \times 2^p}$:

$$M = \begin{pmatrix} \bar{M}_{1,1} & \bar{M}_{1,2} \\ \bar{M}_{2,1} & \bar{M}_{2,2} \end{pmatrix},$$

and let $\bar{M} = (\bar{M}_{i,j})$, a 2×2 block matrix. By order 1 super-symmetry of \bar{M} , M is completely determined by two 2^p dimensional elements in \bar{M} , that is $\bar{M}_{1,1}$ and $\bar{M}_{1,2}$. These elements are themselves order p super-symmetric, hence, they both contain at most 2^p elements. Hence, M contains at most $2 \cdot 2^p$ elements, that is 2^{p+1} elements.

Consider a partition of $[1, 2, \dots, 2^{p+1}]$ into two groups of size 2^p , $P_1 = [1, 2, \dots, 2^p]$ and $P_2 = [2^p + 1, \dots, 2^{p+1}]$. By order 1 super-symmetry of \bar{M} , M is invariant under the permutations $\pi_1^i([P_1, P_2])$, $i = 1, 2$. In addition, by order p super-symmetry of all elements in \bar{M} , there exist 2^p permutations of P_1, P_2 , leaving these elements identical when applied simultaneously. Applying π_1^i to these permutations furnishes 2^{p+1} permutations under which M is invariant. In addition, by construction, and given the fact that none of the elements in P_1 and P_2 are identical, none of these permutations are identical.

Example 8.3. Given $k \in \mathbb{N}$, let $p = 2$, and $M \in (\mathbb{S}^k)^{2^p \times 2^p}$ an order-2 super-symmetric matrix. Then, $M = M_{[1,2,3,4]} = M_{[2,1,4,3]} = M_{[3,4,1,2]} = M_{[4,3,2,1]}$. \square

Building on Definition 8.2, we introduce classes of *super-symmetric* covariance matrices. We then provide, in Section 8.3.5, natural instances of super-symmetric covariance matrices, for which $(P_{\text{relax, multi}})$ becomes equivalent to the initial problem (8.9). Note that the choice of covariance matrices is independent from the noise assumption (8.4).

Definition 8.3 (Super-symmetric covariance class). Let $N, p \in \mathbb{N}$, and $R = 2^p$. Given a class $\Sigma_{N, 2^p}$ of matrices as defined in (8.20), consider the matrix

$$M = \begin{pmatrix} \text{Id}_N & \Sigma_{N, 2^p}^{12} & \Sigma_{N, 2^p}^{13} & \dots & \Sigma_{N, 2^p}^{1R} \\ \Sigma_{N, 2^p}^{12} & \text{Id}_N & \Sigma_{N, 2^p}^{23} & \dots & \Sigma_{N, 2^p}^{2R} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \Sigma_{N, 2^p}^{1R} & \Sigma_{N, 2^p}^{2R} & \Sigma_{N, 2^p}^{3R} & \dots & \text{Id}_N \end{pmatrix}, \quad (8.25)$$

where Id_N is the identity matrix of size N . We say $\Sigma_{N, 2^p}$ is super-symmetric if M is order p super-symmetric.

Relying on symmetry properties of Proposition 8.2, we show that for any $R = 2^p$ ($p \in \mathbb{N}$), $(P_{\text{relax, multi}})$ is strictly equivalent to

$$\begin{aligned} & \max_{\substack{(F, G) \in \mathcal{G}_c \\ G^{1j} \in \mathbb{S}^{3N+4}, j=2, \dots, R}} \text{Tr}(A_{\text{obj}}G) + F^\top b_{\text{opt}}, & (P_{\text{relax, sym}}) \\ & \text{s.t. } (F, G) \text{ is } \mathcal{F}\text{-interpolable, } G \succcurlyeq 0, \\ & \quad G \text{ is order } p \text{ super-symmetric,} \end{aligned}$$

$$\text{where } (\mathbf{F}, \mathbf{G}) = \left(\begin{bmatrix} F \\ F \\ \vdots \\ F \end{bmatrix}, \begin{pmatrix} G & G^{12} & \dots & G^{1R} \\ G^{12} & G & \dots & G^{1(R-1)} \\ \vdots & \vdots & \ddots & \vdots \\ G^{1R} & G^{1(R-1)} & \dots & G \end{pmatrix} \right), \quad (8.26)$$

and $\forall j \in \{1, \dots, R\}, \forall k = 0, 1, \dots, N-1,$

$$\mathbf{v}_k^\top G^{1j} \mathbf{v}_k = (\Sigma_{N,2^p}^{1j})_{k+1,k+1} \mathbf{v}_k^\top G \mathbf{v}_k,$$

where \mathbf{v}_k is the selection vector accessing the element ε_k in G .

Example 8.7: SGD under additive bounded noise model

Let $R = 2$, and consider Example 8.2, i.e., a single iteration of (SGD) on the class $\mathcal{F} = \mathcal{F}_{0,\text{inf}}$ of convex functions, under additive bounded noise. Let $\Sigma_{1,2}^{12} = -1$. Then, $(P_{\text{relax, sym}})$ becomes

$$\max_{\substack{(F,G) \in \mathbb{R}^3 \times \mathbb{S}^6, \\ G^{12} \in \mathbb{S}^6}} F^\top(\mathbf{f}_1 - \mathbf{f}_*) \text{ s.t.} \quad (8.27)$$

$$\begin{pmatrix} G & G^{12} \\ G^{12} & G \end{pmatrix} \succcurlyeq 0$$

$$\begin{aligned} \mathcal{F}\text{-interpolability: } & F^\top(\mathbf{f}_i - \mathbf{f}_j) \geq \mathbf{g}_j^\top G(\mathbf{x}_i - \mathbf{x}_j) \\ & + \frac{1}{2L}(\mathbf{g}_j - \mathbf{g}_i)^\top G(\mathbf{g}_j - \mathbf{g}_i), \quad i, j \in \{0, 1, \star\}, \\ & F^\top(\mathbf{f}_i - \mathbf{f}_j) \geq \mathbf{g}_j^\top G^{12} \mathbf{x}_i - \mathbf{g}_j^\top G \mathbf{x}_j \\ & + \frac{1}{2L} \left(\mathbf{g}_j^\top G \mathbf{g}_j + \mathbf{g}_i^\top G \mathbf{g}_i - 2\mathbf{g}_j^\top G^{12} \mathbf{g}_j \right), \quad i, j \in \{0, 1, \star\}, \end{aligned}$$

$$C_{\text{init}}: \quad (\mathbf{x}_0 - \mathbf{x}_*)^\top G(\mathbf{x}_0 - \mathbf{x}_*) \leq 1,$$

$$C_{\perp}: \quad \mathbf{v}_0^\top G \mathbf{x}_0 = \mathbf{v}_0^\top G \mathbf{g}_0 = 0,$$

$$C_{\text{variance}}: \quad \mathbf{v}_0^\top G \mathbf{v}_0 = \sigma^2,$$

$$\text{Covariance:} \quad \mathbf{v}_0^\top G^{12} \mathbf{v}_0 = -\sigma^2.$$

We now formalize equivalence of $(P_{\text{relax, multi}})$ and $(P_{\text{relax, sym}})$ for super-symmetric covariance classes.

Proposition 8.3 (A symmetric reformulation of $(P_{\text{relax, multi}})$). *Given $N, p \in \mathbb{N}$, let $R = 2^p$, and consider a method \mathcal{M} , a function class \mathcal{F} , a class of distribu-*

tions Ω_N , and an initial condition C_{init} , satisfying Assumption 8.1. In addition, consider an order p super-symmetric covariance class $\Sigma_{N,2^p}$ as defined in (8.20), see Definition 8.3. Let f_{obj}^* be the solution to $(P_{relax, multi})$, and $f_{obj}^{sym,*}$ be the solution to $(P_{relax, sym})$. Then,

$$f_{obj}^{sym,*} = f_{obj}^*.$$

Proof. Clearly, any solution of $(P_{relax, sym})$ is a feasible point of $(P_{relax, multi})$, hence $f_{obj}^{sym,*} \leq f_{obj}^*$. In addition, let

$$(\mathbf{F}^*, \mathbf{G}^*) = \left(\begin{bmatrix} F^1 \\ F^2 \\ \vdots \\ F^R \end{bmatrix}, \begin{pmatrix} G^1 & G^{12} & \dots & G^{1R} \\ G^{12} & G^2 & \dots & G^{2R} \\ \vdots & \vdots & \ddots & \vdots \\ G^{1R} & G^{2R} & \dots & G^R \end{pmatrix} \right)$$

be a solution to $(P_{relax, multi})$.

Claim 1: Given $P = [1, 2, \dots, R]$ and any permutation $\pi_p^i(P)$ as defined in Proposition 8.2,

$$(\mathbf{F}_{\pi_p^i(P)}, \mathbf{G}_{\pi_p^i(P)}),$$

the vector and matrix obtained from (\mathbf{F}, \mathbf{G}) by permuting indices of entries according to $\pi_p^i(P)$, is also a solution to $(P_{relax, multi})$.

Indeed, swapping the order of entries in (\mathbf{F}, \mathbf{G}) leaves the objective function $\frac{1}{R} \sum_{i=1}^R \text{Tr}(A_{obj} G^i) + (F^i)^\top b_{opt}$, the constraint $(F^i, G^i) \in \mathcal{G}_c$, and \mathcal{F} -interpolability of (\mathbf{F}, \mathbf{G}) unchanged. In addition, due to the order p super-symmetry of $\Sigma_{N,R}$, and given Proposition 8.2, the covariance condition is automatically satisfied by $(\mathbf{F}_{\pi_p^i(P)}, \mathbf{G}_{\pi_p^i(P)})$ if satisfied by (\mathbf{F}, \mathbf{G}) .

Claim 2: Given the 2^p permutations $\pi_p^i(P)$ as defined in Proposition 8.2,

$$\begin{aligned} & \frac{1}{R} \sum_{i=1}^R (\mathbf{F}_{\pi_p^i(P)}, \mathbf{G}_{\pi_p^i(P)}) \\ &= \left(\begin{bmatrix} \frac{1}{R} \sum_{i=1}^R F^i \\ \frac{1}{R} \sum_{i=1}^R F^i \\ \vdots \\ \frac{1}{R} \sum_{i=1}^R F^i \end{bmatrix}, \begin{pmatrix} \frac{1}{R} \sum_{i=1}^R G^i & G^{12} & \dots & G^{1R} \\ G^{12} & \frac{1}{R} \sum_{i=1}^R G^i & \dots & G^{1(R-1)} \\ \vdots & \vdots & \ddots & \vdots \\ G^{1R} & G^{1(R-1)} & \dots & \frac{1}{R} \sum_{i=1}^R G^i \end{pmatrix} \right) \end{aligned}$$

is also a solution to $(P_{relax, multi})$. Under Assumption 8.1, all constraints in

$(P_{\text{relax, multi}})$, and its objective constraint, are linear in its variables. Hence, any convex combination of solutions of $(P_{\text{relax, multi}})$ is also a solution to $(P_{\text{relax, multi}})$. In addition, the structure of $\frac{1}{R} \sum_{i=1}^R (\mathbf{F}_{\pi_p^i(P)}, \mathbf{G}_{\pi_p^i(P)})$ follows from the definition of $\pi_p^i(P)$.

Hence, to any solution to $(P_{\text{relax, multi}})$ corresponds a feasible point of $(P_{\text{relax, sym}})$ with same objective value. Therefore, $f_{\text{obj}}^* \leq f_{\text{obj}}^{\text{sym},*}$, which concludes the proof. \square

Letting

$$\tilde{\mathcal{G}}_c(\Sigma_{N, 2^p}) := \left\{ (F, G) \in \mathcal{G}_c : \exists (\mathbf{F}, \mathbf{G}) = \left(\begin{bmatrix} F \\ F \\ \vdots \\ F \end{bmatrix}, \begin{pmatrix} G & G^{12} & \dots & G^{1R} \\ G^{12} & G & \dots & G^{1(R-1)} \\ \vdots & \vdots & \ddots & \vdots \\ G^{1R} & G^{1(R-1)} & \dots & G \end{pmatrix} \right) \right\},$$

s.t. (\mathbf{F}, \mathbf{G}) is \mathcal{F} -interpolable, $\mathbf{G} \succcurlyeq 0$,
 \mathbf{G} is order p super-symmetric,
 and $\forall j \in \{1, \dots, R\}, \forall k = 0, 1, \dots, N-1$,
 $\mathbf{v}_k^\top G^{1j} \mathbf{v}_k = (\Sigma_{N, 2^p}^{1j})_{k+1, k+1} \mathbf{v}_k^\top \mathbf{G} \mathbf{v}_k$

it holds that $\mathcal{G}_E \subseteq \tilde{\mathcal{G}}_c(\Sigma_{N, R}) \subseteq \mathcal{G}_c$, and

$$\begin{aligned} & \max_{(F, G) \in \mathcal{G}_E} \text{Tr}(A_{\text{obj}} G) + F^\top b_{\text{opt}} \\ & \leq \max_{(F, G) \in \tilde{\mathcal{G}}_c(\Sigma_{N, R})} \text{Tr}(A_{\text{obj}} G) + F^\top b_{\text{opt}} \\ & \leq \max_{(F, G) \in \mathcal{G}_c} \text{Tr}(A_{\text{obj}} G) + F^\top b_{\text{opt}}. \end{aligned}$$

The sets $\tilde{\mathcal{G}}_c(\Sigma_{N, R})$ thus yields, for different choices of $\Sigma_{N, R}$, a spectrum of refined SDP relaxations of (8.9), starting from $(P_{\text{relax, 1}})$ and, as will be proven in Section 8.3.5, eventually attaining (8.9).

Computational comments When solving $(P_{\text{relax, sym}})$, storing the full matrix \mathbf{G} is unnecessary. The \mathcal{F} -interpolability constraint only involves submatrices G, G^{1j} , and the constraint $\mathbf{G} \succcurlyeq 0$ can be equivalently reduced to linear inequalities on G, G^{1j} , as shown in Appendix 8.A.

8.3.5 Exact SDP reformulation of (8.9)

We now finally propose a SDP equivalent reformulation of (8.9), by introducing a covariance matrix $\Sigma_{N,R}$ for which $(P_{\text{relax, sym}})$ is equivalent to (8.9). That is, setting $R = 2^N$, we introduce $\Sigma_{N,2^N}^*$, for which

$$\max_{(F,G) \in \mathcal{G}_E} \text{Tr}(A_{\text{obj}}G) + F^\top b_{\text{opt}} = \max_{(F,G) \in \tilde{\mathcal{G}}_c(\Sigma_{N,2^N}^*)} \text{Tr}(A_{\text{obj}}G) + F^\top b_{\text{opt}}.$$

This equivalence holds regardless of the noise variance (8.4), since $\Sigma_{N,R}$ is independent of this assumption. The covariance matrix ensures that the copies in G_c correspond to scenarios generated by a noise of equiprobable 2-realizations distribution.

Definition 8.4 (Completely anti-correlated covariance class). Let $N \in \mathbb{N}$, and $R = 2^N$. We say $\Sigma_{N,2^N}^*$ as defined in (8.20) is *completely anti-correlated* if (i) $\Sigma_{N,2^N}^*$ is order N super-symmetric, and (ii) the generating matrices Σ^{1j} , for $j = 2, \dots, 2^N$, have diagonals given by all vectors in $\{-1, 1\}^N$ ordered lexicographically, excluding the identity vector $\mathbf{1}_N$ (that would correspond to Σ^{11}), and enumerated in reverse external order, that is, starting from the closest vector to $\mathbf{1}_N$.

Example 8.4. Let $N = 2$. Then, $\Sigma_{N,R}$ is completely anti correlated if order 2 super-symmetric and defined via

$$\Sigma^{12} = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}, \Sigma^{13} = \begin{pmatrix} -1 & 0 \\ 0 & 1 \end{pmatrix}, \Sigma^{14} = \begin{pmatrix} -1 & 0 \\ 0 & -1 \end{pmatrix}.$$

Theorem 8.3 (Exact SDP formulation of (8.9))

Given $N \in \mathbb{N}$, let $R = 2^N$, and consider a method \mathcal{M} , a function class \mathcal{F} , a class of distributions Ω_N , and an initial condition C_{init} , satisfying Assumption 8.1. In addition, consider a completely anti-correlated covariance class $\Sigma_{N,2^N}^*$ as defined in (8.20), see Definition 8.3. Then,

$$\max_{(F,G) \in \mathcal{G}_E} \text{Tr}(A_{\text{obj}}G) + F^\top b_{\text{opt}} = \max_{(F,G) \in \tilde{\mathcal{G}}_c(\Sigma_{N,R})} \text{Tr}(A_{\text{obj}}G) + F^\top b_{\text{opt}}.$$

Proof. By Propositions 8.2 and 8.3,

$$\max_{(F,G) \in \mathcal{G}_E} \text{Tr}(A_{\text{obj}}G) + F^\top b_{\text{opt}} \leq \max_{(F,G) \in \tilde{\mathcal{G}}_c(\Sigma_{N,R})} \text{Tr}(A_{\text{obj}}G) + F^\top b_{\text{opt}}.$$

We now show that any $(F, G) \in \tilde{\mathcal{G}}_c(\Sigma_{N,R})$ belongs to \mathcal{G}_E .

Denote by $(F^{(i)}, G^{(i)})$ any of the 2^N copies of (F, G) . We argue that $(F^{(i)}, G^{(i)})$ are the realizations of a random variable, generated by realizations $(\varepsilon^{(i)}, \sigma^2)$ of $(\varepsilon, \sigma^2) \in \Omega_N$, where σ^2 is deterministic and ε is an equidistributed random variable in Ω_N , with, at each step, 2 possible realizations, i.e.,

$$\begin{aligned} \varepsilon_k^{(i)} &= \pm v_k \sqrt{A_1(f(x_k) - f_\star) + B_1\|g_k\|^2 + C_1\|x_k - x_\star\|^2 + D_1 + E_1\sigma_k^2} \\ \sigma_{k+1}^2 &= A_2(f(x_k) - f_\star) + B_2\|g_k\|^2 + C_2\|x_k - x_\star\|^2 + D_2 + (1 - \rho)\sigma_k^2, \end{aligned}$$

where $v_k \in \mathbb{R}^d$ is a vector of norm 1, orthogonal to all past information at iteration k , and may differ from query to query. The 2^N realizations $\varepsilon^{(i)}$ are generated by all vectors in $\{-1, 1\}^N$, ordered lexicographically. Indeed, given $\Sigma_{N,2^N}^\star$, the constraints $\mathbf{v}_k^\top G^{1i} \mathbf{v}_k = \sum_{k+1, k+1}^{1i} \mathbf{v}_k^\top G \mathbf{v}_k$ in $\tilde{\mathcal{G}}_c(\Sigma_{N,2^N}^\star)$, are satisfied by such realizations $\varepsilon^{(i)}$.

In addition, by \mathcal{F} -interpolability of (\mathbf{F}, \mathbf{G}) , there exists a single function $f \in \mathcal{F}$ interpolating all $(F^{(i)}, G^{(i)})$, and these are all consistent with the method \mathcal{M} , the optimality condition, and C_{init} by definition of $\tilde{\mathcal{G}}_c(\Sigma_{N,2^N}^\star)$. Hence, $(F, G) = \frac{1}{2^N} \sum_i (F^{(i)}, G^{(i)})$ belongs to \mathcal{G}_E . \square

Theorem 8.3 shows that the worst-case instance in \mathcal{G}_E corresponds to noise with a very simple equiprobable 2-realizations distribution, yielding identical realizations $(F^{(i)}, G^{(i)})$. This aligns with intuition, as the objective depends symmetrically on all realizations, and was already showcased in [DS20, Theorems 2, 3]. Building on Theorem 8.3, we finally show that the initial set \mathcal{G}_E , as defined in (8.14), is convex.

Theorem 8.4

Given $N \in \mathbb{N}$, consider a method \mathcal{M} , a function class \mathcal{F} , a class of distributions Ω_N , and an initial condition C_{init} , satisfying Assumption 8.1. Then, the set \mathcal{G}_E as defined in (8.14) is convex.

Proof. The proof of Theorem 8.2 showed that $\mathcal{G}_E \subseteq \tilde{\mathcal{G}}_c(\Sigma_{N,2^N}^\star)$. In addition, the proof of Theorem 8.3 showed that $\tilde{\mathcal{G}}_c(\Sigma_{N,2^N}^\star) \subseteq \mathcal{G}_E$. Hence, $\mathcal{G}_E =$

$\tilde{\mathcal{G}}_c(\Sigma_{N,2N}^*)$. Convexity of \mathcal{G}_E follows from that of $\tilde{\mathcal{G}}_c(\Sigma_{N,2N}^*)$, which can be expressed as the projection of a convex set, since characterized by linear constraints. \square

Hence, (8.9) can be equivalently formulated as (i) a maximization problem over a convex set of size $\propto 2^N$, as in Theorem 8.3, or (ii) a maximization problem over a convex set \mathcal{G}_E of size $\propto N$, as in (P_{exact}) . In addition, \mathcal{G}_E is the intersection of a set determined by interpolation and noise model constraints, independent of algorithmic specifications; and linear constraints encoding the algorithm. These linear constraints can thus be interpreted as additional constraints in the maximization problem. This formulation is similar to the deterministic setting, see, e.g., [THG17c]. We were however unable to obtain a tractable formulation of this set of size $\propto N$, hence obtaining a tight reformulation of (8.9) requires, for the moment, solving a SDP of size $\propto 2^N$. But, as shown in Section 8.4, in many cases, the relaxation $(P_{\text{relax},1})$, i.e., an SDP of size $\propto N$, is sufficient to obtain tight convergent rates.

8.4 Applications: performance analysis of SGD on $\mathcal{F}_{\mu,L}$

This section relies on the framework introduced in Section 8.3 to analyze (SGD) on $\mathcal{F}_{\mu,L}$, under several noise models. We compare the resulting convergence rates with existing guarantees, obtained classically or in a computer-aided way.

Given, as only information on the noise distribution, its expectation, variance, and uncorrelation conditions, the tightest formulation of the proposed framework provides the best rates possible. We demonstrate both numerically (Figures 8.1, 8.3, 8.4 and 8.5) and analytically (Theorem 8.5 and Proposition 8.4) how it improves on existing bounds derived on the basis of this information only.

However, as compared to the all-scenarios approach introduced in [TB19], which uses, when available, more information on the distribution of the noise than its expectation, variance, and uncorrelation conditions, an information that our framework cannot use, and hence applies to a more restrictive set of problems, the proposed framework may yield worse guarantees, as it is typically the case for linear convergence rates (Figures 8.3, 8.5). On the other hand, the all-scenarios approach is restricted to the anal-

ysis of finite-support noise settings and, mainly, to a potential function approach. In addition, it yields better rates only if further information on the noise distribution is available, e.g., in the finite-sum and the block coordinate settings. Hence, on (i) noise models defined only via their expectation, variance, and uncorrelation, e.g., the absolute bounded model (Figure 8.1), or (ii) problems with sublinear convergence rates (Figure 8.4b), the proposed framework improves on existing results yielded by the all-scenarios approach.

In addition, the framework significantly improves on rates obtained via the worst-scenario approach, i.e., under deterministic perturbation (Figures 8.1 and 8.1). This demonstrates the necessity, in situations where the perturbation is in fact stochastic rather than deterministic, to analyze methods under a stochastic framework.

Finally, we observe that, in many settings, solving $(P_{\text{relax},1})$ (relaxing \mathcal{G}_E by \mathcal{G}_c), called the single realization approach, already yields exact convergence rates, since it matches the exact approach $(P_{\text{relax},\text{sym}})$ with $R = 2^N$ and $\Sigma_{N,2^N}^*$ (cf. Definition 8.4). In what follows, we refer to $(P_{\text{relax},\text{sym}})$ with $R = 2$ and $\Sigma_{N,R} = \mathbf{1}_N$, the all-ones vector of size N (relaxing \mathcal{G}_E by $\tilde{\mathcal{G}}_c(-\mathbf{1}_N)$), as the two-realizations approach; and to $(P_{\text{relax},\text{sym}})$ with $R = 2^N$ and $\Sigma_{N,2^N}^*$, as the 2^N -realizations exact approach.

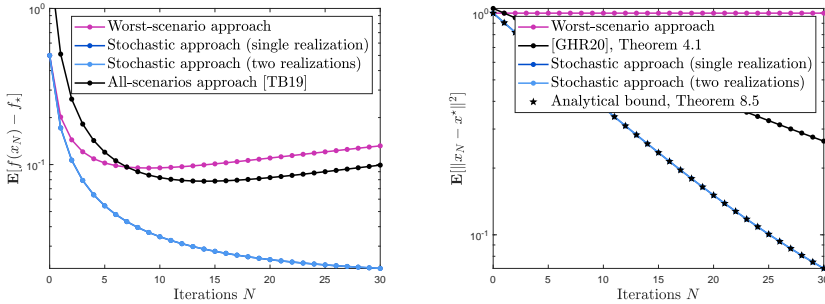
Additive bounded noise model. Consider (SGD) with constant stepsizes α , under the additive bounded noise model, i.e., $\mathbb{E}[\|\varepsilon_k\|^2] \leq \sigma^2$. Figure 8.1 compares various convergence rates, in both the convex and the strongly convex regimes. Specifically, it displays the worst-case scenario, or equivalently the analysis of (SGD) under a deterministic bounded perturbation; the PEP-based bounds obtained via the single realization and two-realizations approaches; and, in the convex case, the all-scenarios PEP-based bound of [TB19, Theorem 5] for $\alpha = \frac{1}{L}$, that is

$$\mathbb{E}[f(x_N) - f_\star] \leq \frac{L}{2N} \|x_0 - x_\star\|^2 + \frac{N+3}{4L} \sigma^2. \quad (8.28)$$

In addition, in the strongly convex case, it displays an existing bound from [GHR20, Theorem 4.1], that is,

$$\mathbb{E}[\|x_N - x_\star\|^2] \leq \phi \|x_0 - x_\star\|^2 + \frac{\sigma^2 \alpha}{\mu}, \text{ where } \phi = 1 - \mu \alpha. \quad (8.29)$$

This bound is derived based on the weaker assumption $E[\|\varepsilon_k\|^2 | \mathcal{A}_k] \leq 2L(f(x_k) - f(x_*) + \sigma^2)$, satisfied by the additive bounded noise setting, see, e.g., [Sti19, Example 2].



(a) Computer-aided convergence analyses of (SGD) with constant stepsizes $\alpha = \frac{1}{L}$ on $\mathcal{F}_{0,L}$ under the additive bounded noise model, with performance measure $\mathbb{E}[f(x_N) - f_*$], and initial condition $\|x_0 - x_*\| \leq 1$, for $L = 1$.

(b) Convergence analyses of (SGD) with constant stepsizes $\alpha = \frac{1}{2L}$ on $\mathcal{F}_{\mu,L}$, $\mu = 0.1$, under the additive bounded noise model, with performance measure $\mathbb{E}[\|x_N - x_*\|^2]$, and initial condition $\|x_0 - x_*\| \leq 1$, for $L = 1$.

Fig. 8.1 Comparison of computer-aided convergence analyses of (SGD) with constant stepsizes α , on $\mathcal{F}_{0,L}$ (convex L -smooth functions, left-hand side figure), and $\mathcal{F}_{\mu,L}$ (μ -strongly convex L -smooth functions, $\mu > 0$, right-hand side figure), under the additive bounded noise model of variance $\sigma^2 = 0.01$. Displayed are the worst-scenario PEP-based bound (pink), the PEP-based bound from the framework presented in Section 8.3 (blue), resulting from the single realization and the 2-realizations approaches, the all scenarios PEP-based bound from [TB19](black), the bound from [GHR20, Theorem 4.1](black), and the analytical bound from Theorem 8.5 (\star).

The single and two realizations approaches yield identical bounds, indicating that enforcing a common interpolating function across two noise realizations does not improve the guarantee. The bound yielded by the 2^N -realizations exact approach, computed for $N \leq 6$ (not displayed for the sake of clarity), turns out to also be equivalent to these bounds, showing that the bound resulting from $(P_{\text{relax},1})$ is, in fact, tight for this noise model, up to $N = 6$. Moreover, this bound outperforms [GHR20, Theorem 4.1] and [TB19, Theorem 5]. Although [TB19, Theorem 5] derives from an all-scenarios formulation, it considers as performance measure a slightly different setting, that is, $\mathbb{E}[f(x_N) - f(x_*) + \frac{L}{2N}\|x_N - x_*\|^2]$, from which we infer a bound on $\mathbb{E}[f(x_N) - f(x_*)]$. This explains why it is initially outper-

formed by the worst-scenario guarantee.

Finally, in the strongly convex case, relying on numerical insight, we identify an analytical convergence rate that exactly matches the numerical bound from $(P_{\text{relax},1})$, and prove its tightness.

Theorem 8.5. *Let $N \in \mathbb{N}$, $0 < \mu < L$, $\sigma^2 \in \mathbb{R}$, $\alpha \in (0, \frac{2}{L+\mu}]$. Consider N iterations of (SGD) with constant stepsize α , applied to a function $f \in \mathcal{F}_{\mu,L}$, under additive bounded noise $(\varepsilon_k)_{k=0,1,\dots,N-1}$ satisfying $\mathbb{E}[\varepsilon_k | \mathcal{A}_k] = 0$, and $\mathbb{E}[\varepsilon_k | \mathcal{A}_k]^2 \leq \sigma^2$. Let $(x_k)_{k=0,1,\dots,N}$ be the iterates of (SGD), and x_* denote a minimizer of f . Then,*

$$E[\|x_N - x_*\|^2] \leq \phi^{2N} \|x_0 - x_*\|^2 + \frac{1 - \phi^{2N}}{1 - \phi^2} \alpha^2 \sigma^2, \quad (8.30)$$

where $\phi = (1 - \mu\alpha)$. In addition, this bound is tight.

Proof. The first part of the proof follows exactly [Tay24, Section 2.8.1], which addresses the noise-free case. Let $S = \{(x_i, f_i, g_i)\}_{i=0,1,\dots,N,*}$ denote the set of iterates (and minimizer) and associated function and gradient values. We start by analyzing a single iteration of (SGD), and claim the following.

$$\|x_{k+1} - x_*\|^2 \leq \phi^2 \|x_k - x_*\|^2 + \alpha^2 \sigma^2. \quad (8.31)$$

By Proposition 2.1 describing smooth strongly convex functions, and since $g_* = 0$, it holds

$$\begin{aligned} f_* &\geq f_k + \langle g_k, x_* - x_k \rangle \\ &\quad + \frac{L}{2(L-\mu)} \left(\frac{\|g_k\|^2}{L} + \mu \|x_k - x_*\|^2 - \frac{2\mu}{L} \langle g_k, x_k - x_* \rangle \right). \\ f_k &\geq f_* + \frac{L}{2(L-\mu)} \left(\frac{\|g_k\|^2}{L} + \mu \|x_k - x_*\|^2 - \frac{2\mu}{L} \langle g_k, x_k - x_* \rangle \right). \end{aligned}$$

Summing these inequalities, multiplied by $2\alpha\phi$, yields

$$-\frac{2\alpha\phi(L+\mu)}{L-\mu} \langle g_k, x_k - x_* \rangle + \frac{2\alpha\phi}{L-\mu} \|g_k\|^2 + \frac{2\alpha\phi\mu L}{L-\mu} \|x_k - x_*\|^2 \leq 0,$$

hence

$$\|x_k - x_*\|^2 - 2\alpha \langle g_k, x_k - x_* \rangle + \alpha^2 \|g_k\|^2 \leq \phi^2 \|x_k - x_*\|^2$$

$$\begin{aligned}
 & - \frac{\alpha(2 - \alpha(L + \mu))}{L - \mu} \|\mu(x_k - x_*) - g_k\|^2 \\
 & \leq \phi^2 \|x_k - x_*\|^2,
 \end{aligned}$$

where the second inequality follows from the domain of α . Hence,

$$\begin{aligned}
 & \|x_k - x_*\|^2 - 2\alpha \langle x_k - x_*, g_k + \varepsilon_k \rangle + \alpha^2 \|g_k + \varepsilon_k\|^2 \leq \phi^2 \|x_k - x_*\|^2 \\
 & \quad - 2\alpha \langle x_k - x_*, \varepsilon_k \rangle + \alpha^2 (2\langle g_k, \varepsilon_k \rangle + \|\varepsilon_k\|^2) \\
 \Leftrightarrow & \|x_{k+1} - x_*\|^2 \leq \phi^2 \|x_k - x_*\|^2 - 2\alpha \langle x_k - x_*, \varepsilon_k \rangle + \alpha^2 (2\langle g_k, \varepsilon_k \rangle + \|\varepsilon_k\|^2).
 \end{aligned}$$

Taking the expectation with respect to ε_k (conditioned on x_k) yields

$$\begin{aligned}
 \mathbb{E}_{\varepsilon_k} [\|x_{k+1} - x_*\|^2] & \leq \phi^2 \|x_k - x_*\|^2 + \alpha^2 \mathbb{E}_{\varepsilon_k} [\|\varepsilon_k\|^2] \\
 & \leq \phi^2 \|x_k - x_*\|^2 + \alpha^2 \sigma^2,
 \end{aligned}$$

by assumption on ε_k . Hence, (8.31) holds. By recursion, it holds

$$\mathbb{E} [\|x_N - x_*\|^2] \leq \phi^{2N} \|x_0 - x_*\|^2 + \sum_{i=0}^{N-1} \phi^{2i} \alpha^2 \sigma^2.$$

Since, for $r \neq 1$, $\sum_{i=0}^{N-1} r^i = \frac{1-r^N}{1-r}$, this concludes the derivation of the bound.

To conclude the proof, observe that the function $f(x) = \frac{\mu}{2} \|x\|^2$ and any distribution satisfying $\mathbb{E}[\varepsilon_k] = \sigma^2$, $k = 0, 1, \dots, N$, attain the bound. Indeed, then, $x_* = 0$, and

$$x_{k+1} = \phi x_k - \alpha \varepsilon_k.$$

Hence,

$$\begin{aligned}
 (x_{k+1} - x_*)^2 & = (\phi(x_k - x_*) - \alpha \varepsilon_k)^2 = \phi^2 (x_k - x_*)^2 - 2\alpha \phi \varepsilon_k (x_k - x_*) + \alpha^2 \varepsilon_k^2 \\
 \Leftrightarrow \mathbb{E}[(x_{k+1} - x_*)^2 | \mathcal{A}_k] & = \phi^2 (x_k - x_*)^2 + \alpha^2 \sigma^2.
 \end{aligned}$$

□

Relative additive noise and the silver stepsize schedule: on the gap between \mathcal{G}_c and \mathcal{G}_E . Recently, [AP24] proposed a stepsize schedule accelerating gradient descent ((SGD), $\varepsilon_k = 0$) on $\mathcal{F}_{0,L}$ in the deterministic setting.

For $N = 2^k - 1, k \in \mathbb{N}$, the schedule is defined as

$$\alpha = \begin{cases} \sqrt{2}, & \text{if } k = 1, \\ \alpha = [\alpha, 1 + (1 + \sqrt{2})^{j-1}, \alpha], & j = 1, \dots, k - 1, \text{ otherwise.} \end{cases} \quad (8.32)$$

As shown in [VG24] via a worst-scenario PEP analysis, this schedule is highly sensitive to relative additive perturbation. Figure 8.2 demonstrates that this sensitivity significantly diminishes when considering stochastic noise rather than deterministic perturbations.




Fig. 8.2 Comparison of computer-aided convergence analyses of (SGD) on $\mathcal{F}_{\mu,L}$ with relative additive noise ($\mathbb{E}[\varepsilon_k | \mathcal{A}_k] = 0, \mathbb{E}[\|\varepsilon_k | \mathcal{A}_k\|^2] \leq \sigma^2 \|g_k\|^2$). We use the performance measure $\mathbb{E}[\|x_N - x_*\|^2]$, initial condition $\|x_0 - x_*\| \leq 1, \sigma = 0.15, \mu = 0.1, L = 1$, and stepsizes as in (8.32). Displayed are the PEP-based bound treating noise as deterministic (pink), the stochastic noise bounds (blue), resulting from the single realization approach, the two realizations approach, and the exact 2^N realizations approach (up to $N = 3$), as well as the PEP-based bound in the exact setting (black).

Remark 8.2 (Symmetric solutions). In Figure 8.1, we noted that the bound from $(P_{\text{relax},1})$ coincided with that of (P_{exact}) , for $N \leq 6$, implying that the solution to $(P_{\text{relax},1})$ actually lies in \mathcal{G}_E , up to at least $N \leq 6$. This occurs whenever the problem's solution is sufficiently symmetric, intuitively, when the worst-case function with respect to any noise realization remains the worst-case function for any other realization.

Many problems exhibit this symmetry. However, for schedules like

the silver stepsize or other stepsizes with occasional longer steps, see, e.g., [Gri24], this symmetry breaks. Accordingly, Figure 8.2 shows that the two-realizations approach yields a modest improvement over the single realization approach. Up to $N = 3$, the 2^N -realizations exact approach yields the same rate as the 2-realizations approach.

In what follows, the single realization approach is always equivalent to the two realizations approach and, up to $N = 6$, to the exact 2^N realizations approach. Hence, we refer to the "stochastic noise" approach as the single realization approach, which corresponds to the exact approach up to $N = 6$.

Finite-sum setting: on the importance of the parameters in (8.4). We now turn to the analysis of (SGD) with constant stepsizes, under noise satisfying a variance condition of the form

$$\mathbb{E}[\|\varepsilon_k\|^2 | \mathcal{A}_k] \leq A_1(f(x_k) - f(x_*)) + D_1. \quad (8.33)$$

This includes, in particular, the finite-sum setting where $f = \sum_i f_i$ is a sum of convex L_i -smooth functions with $L = \max_i L_i$, and the variance at the optimum is bounded, i.e., $\mathbb{E}[\|\nabla f_i(x_*)\|^2] := \sigma_*^2 < +\infty$. Indeed, as shown in, e.g., [GG23, Lemma 4.20], this setup satisfies (8.33) with $A_1 = 4L$ and $D_1 = 2\sigma_*^2$. Moreover, when $\sigma_*^2 = 0$, e.g., in overparametrized models [MBB18], tighter constants can be used: $A_1 = 2L$, $D_1 = 0$; see [GG23, Lemma 4.19].

In [GHR20, Theorem 4.1], a convergence rate is derived on the basis of (8.33), that is, using only the expectation, variance, and uncorrelation properties of the noise. This bound is given by

$$E[\|x_N - x_*\|^2] \leq \phi^N \|x_0 - x_*\|^2 + \frac{\alpha^2 D_1}{\alpha \mu}, \text{ where } \phi = (1 - \mu\alpha). \quad (8.34)$$

On the other hand, the finite-sum setting was analyzed using the exact all-scenarios approach in [TB19] and [CKPG25], using all information available on the noise distribution (hence, for a more restrictive setting than $(P_{\text{relax},1})$). In particular, [CKPG25, Theorem 3.3] shows

$$E[\|x_N - x_*\|^2] \leq \phi^N \|x_0 - x_*\|^2 + \frac{1 - \phi^N}{1 - \phi} \frac{\alpha^2 D_1}{2} \left(1 + \frac{\alpha(L - \mu)}{2 - (L + \mu)\alpha} \right), \quad (8.35)$$

where $\phi = (1 - \mu\alpha)^2$. Figure 8.3 compares the proposed framework with

these existing rates from [GHR20, Theorem 4.1] and [CKPG25, Theorem 3.3]. The right-hand side of Figure 8.3 illustrates the impact of the parameter choice in the variance model. For instance, in the case $\sigma_\star^2 = 0$, using $A_1 = 4L$ (as in the general case) rather than the sharper $A_1 = 2L$ yields looser performance guarantees.

The bound [GHR20, Theorem 4.1] is obtained by (i) showing satisfaction, for all $k = 0, 1, \dots, N - 1$, of [GHR20, Lemma C1]:

$$E[\|x_{k+1} - x_\star\|^2] \leq \phi \|x_k - x_\star\|^2 + \alpha^2 D_1, \text{ where } \phi = (1 - \mu\alpha), \quad (8.36)$$

and (ii) recursively applying (8.36). In the computations, $\sum_{i=0}^{N-1} \phi^i$ is relaxed into $\frac{1}{1-\phi}$. We strengthen [GHR20, Theorem 4.1] by removing this relaxation.

Proposition 8.4. *Let $N \in \mathbb{N}$, $0 < \mu < L$, $A_1, D_1 \in \mathbb{R}$, $\alpha \in (0, \frac{2}{A_1}]$. Given a function $f \in \mathcal{F}_{\mu,L}$, denote by x_\star one of its minimizers. Consider N iterations of (SGD) with constant stepsize $\alpha_k = \alpha$, $k = 0, 1, \dots, N$, on f , with additive noise $(\varepsilon_k)_{k=0,1,\dots,N-1}$ satisfying $\mathbb{E}[\varepsilon_k | \mathcal{A}_k] = 0$, and $\mathbb{E}[\|\varepsilon_k\|^2 | \mathcal{A}_k] \leq A_1(f(x_k) - f(x_\star)) + D_1$, where $(x_k)_{k=0,1,\dots,N}$ denote the iterates of (SGD). Then,*

$$E[\|x_N - x_\star\|^2] \leq \phi^N \|x_0 - x_\star\|^2 + \frac{1 - \phi^N}{1 - \phi} \alpha^2 D_1, \quad (8.37)$$

where $\phi = (1 - \mu\alpha)$.

Proof. It suffices to perform a recursion on (8.36), and rely on $\sum_{i=0}^{N-1} \phi^i = \frac{1 - \phi^N}{1 - \phi}$. \square

The analytical guarantee in Proposition 8.4 coincides with the numerical bound from $(P_{\text{relax},1})$ (exact up to at least $N \leq 6$), for $\alpha = \frac{2}{A_1}$. Additionally, the bound from [GHR20, Theorem 4.1] matches exactly when $D_1 = 0$. As expected, the exact bound under the finite-sum setting is tighter than those based solely on expectation and variance of the noise. In addition, multiple parameter choices for the noise variance are possible in the finite-sum setting, resulting in bounds whose quality depends on these parameters. However, the bounds from [GHR20, Theorem 4.1], $(P_{\text{relax},1})$, and Proposition 8.4, cover more stochastic settings than finite-sum problems.

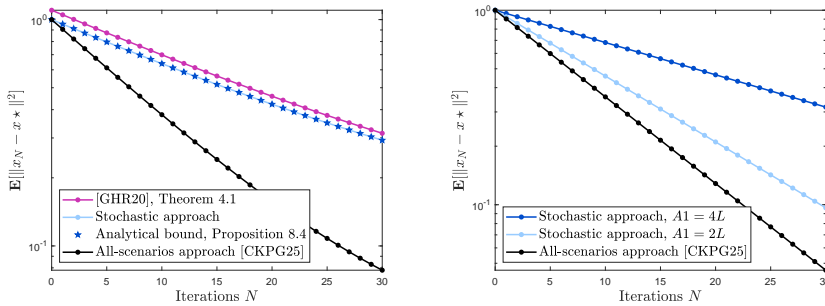


Fig. 8.3 Comparison of convergence analyses of (SGD) with constant stepsizes $\alpha = \frac{1}{2L}$, on $\mathcal{F}_{\mu,L}$, under noise satisfying $\mathbb{E}[\varepsilon_k | \mathcal{A}_k] = 0$, $\mathbb{E}[\|\varepsilon_k\|^2 | \mathcal{A}_k] \leq A_1(f(x_k) - f_*) + D_1$. The performance criterion is $\mathbb{E}[\|x_N - x_*\|^2]$, with initial condition $\|x_0 - x_*\| \leq 1$, and parameters $\mu = 0.1$, $L = 1$, and $D_1 = 0.01$ (left-hand side figure) or $D_1 = 0$ (right-hand side figure). Displayed are the bound from [GHR20, Theorem 4.1], based solely on the expectation, variance, and uncorrelation of the noise (pink), the PEP-based bounds from $(P_{\text{relax},1})$, exact up to $N = 6$ (blue), the analytical bound from Proposition 8.4 (\star), and the exact all-scenarios PEP-based bound from [CKPG25, Theorem 3.3], considering the more restrictive setting of finite-sum optimization (black).

Strong growth condition and RCD We now consider noise satisfying the strong growth condition, i.e., for some $d \in \mathbb{N}$,

$$\mathbb{E}[\|\varepsilon_k\|^2 | \mathcal{A}_k] \leq (d - 1)\|g_k\|^2,$$

introduced in [SR13]. This condition is satisfied by several settings, including some finite-sum problems, where it implies $\mathbb{E}[\|\varepsilon_k\|^2 | \mathcal{A}_k] \leq 2Ld(f(x_k) - f(x_*))$, see [VBS19, Proposition 1], and randomized coordinate descent (RCD), where the gradient estimator is $d\nabla_i f(x_k)\mathbf{e}_i$, with d the dimension and \mathbf{e}_i a randomly selected basis vector [Wri15, Section 1.4]. This estimator is chosen as to be zero-mean, and the classical estimator $\nabla_i f(x_k)\mathbf{e}_i$ can be analyzed via a modified stepsize, $\tilde{\alpha} = d\alpha$. Throughout, we set $\alpha = \frac{1}{d}$.

Figure 8.4a compares, in the strongly convex regime, the bound resulting from $(P_{\text{relax},1})$, exact up to $N = 6$, the bound from [GHR20, Theorem 4.1] (based on the variance model $\mathbb{E}[\|\varepsilon_k\|^2 | x_k] \leq 2Ld(f(x_k) - f(x_*))$), and the exact all-scenarios approach from [TB19, Theorem 17], specifically adapted

to RCD. This bound is given by

$$\mathbb{E}[\|x_N - x_\star\|^2] \leq \left(\frac{(1 - \frac{\mu}{L})^2 + d - 1}{d} \right)^N \|x_0 - x_\star\|^2. \quad (8.38)$$

The bound from [AdKZ23b], solving a slightly relaxed version of $(P_{\text{relax},1})$, exactly matches $(P_{\text{relax},1})$.

Figure 8.4b compares, in the convex regime, the bound resulting from the proposed framework, exact up to $N = 6$, the bound resulting from the approach proposed in [AdKZ23b], and the all-scenarios approach from [TB19, Theorem 17], i.e., letting $\tilde{\alpha} = \frac{1}{L}$,

$$\mathbb{E}[f(x_N) - f_\star] \leq \frac{d\tilde{\alpha}}{2N} \|x_0 - x_\star\|^2. \quad (8.39)$$

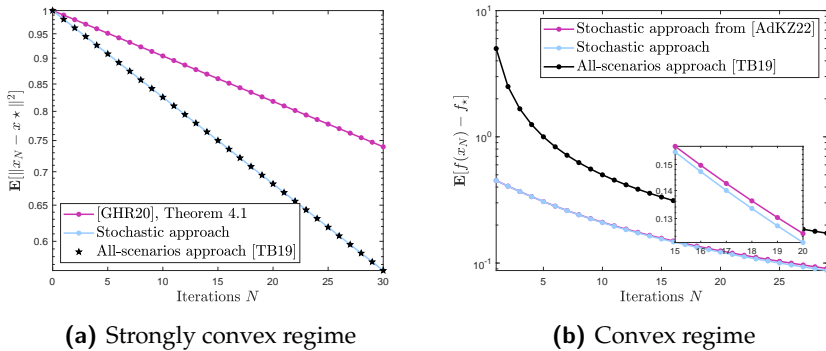


Fig. 8.4 Comparison of convergence analysis of (SGD) with constant stepsize $\alpha_k = \frac{1}{dL}$, on $\mathcal{F}_{\mu,L}$ ($\mu > 0$) and $\mathcal{F}_{0,L}$, with noise satisfying $\mathbb{E}[\varepsilon_k | \mathcal{A}_k] = 0$, and $\mathbb{E}[\|\varepsilon_k\|^2 | \mathcal{A}_k] \leq (d - 1)\|g_k\|^2$. The performance measure is $\mathbb{E}[\|x_N - x_\star\|^2]$ with initial condition $\|x_0 - x_\star\| \leq 1$, for $\mu = 0.1$, $L = 1$, and $d = 10$. The left-hand side figure, considering the strongly convex regime, displays the bound from [GHR20, Theorem 4.1], using $\mathbb{E}[\|\varepsilon_k\|^2 | x_k] \leq 2Ld(f(x_k) - f(x_\star))$ (pink), the PEP-based bound from $(P_{\text{relax},1})$ (blue), and the exact all-scenarios bound designed for RCD from [TB19, Theorem 17] (black). The right-hand side figure, considering the convex regime, displays the PEP-based bound from $(P_{\text{relax},1})$, exact up to $N = 6$ (blue), the PEP-based bound from [AdKZ23b] (pink), and the all-scenarios bound designed for RCD from [TB19, Theorem 16] (black).

Surprisingly, the bounds resulting from the all-scenario approach, i.e.,

[TB19, Theorems 16, 17], and that rely on more information than the framework of Section 8.3, (i) exactly match the bounds resulting from $(P_{\text{relax},1})$ in the strongly convex setting, and (ii) are outperformed by $(P_{\text{relax},1})$ in the convex setting. This lack of tightness is explained by the structure of the potential function, chosen by hand in the all-scenario approach. In addition, the bound from $(P_{\text{relax},1})$ improves on [GHR20, Theorem 4.1], and, in the convex setting (sublinear convergence rate), on the one of [AdKZ23b].

Variance-reduction method: SAGA Finally, we consider the variance-reduction method SAGA [DBLJ14], designed for finite-sum optimization problems. SAGA update at iteration k is

$$x_{k+1} = x_k - \alpha \left(\nabla f_{i_k}(x_k) - y_{i_k}^k + \frac{1}{n} \sum_{j=1}^n y_j^k \right),$$

where i_k is sampled uniformly from $\{1, \dots, n\}$, and y_j^k stores the most recent gradient ∇f_j evaluated at some previous iterate. SAGA can be embedded as (SGD), with constant stepsizes α , and $\nabla f_{i_k}(x_k) - y_{i_k}^k + \frac{1}{n} \sum_{j=1}^n y_j^k$ an unbiased estimator of $\nabla f(x_k)$, satisfying $\mathbb{E}[\varepsilon_k | \mathcal{A}_k] = 0$, and the variance condition (8.4) with parameters

$$A_1 = 4L, \quad B_1 = -1, \quad E_1 = 2, \quad \rho = \frac{1}{n}, \quad A_2 = \frac{2L}{n},$$

where $L = \max_i L_i$, n is the number of component functions, and $\sigma_k^2 = \frac{1}{n} \sum_{i=1}^n \|y_i^k - \nabla f_i(x_*)\|^2$ [GHR20, Lemma A.6].

Figure 8.5 compares the convergence bounds from [GHR20, Theorem 4.1] (based only on noise expectation and variance), the bound from $(P_{\text{relax},1})$, and the exact all-scenarios bound after one iteration, from [GMG⁺22].

The exact all-scenarios bound outperforms $(P_{\text{relax},1})$ but suffers from scalability issues. In addition, it depends on n , making analytical characterization difficult. The bound from [GHR20, Theorem 4.1] coincides with $(P_{\text{relax},1})$ after one iteration, but loses its tightness property from 2 iterations.

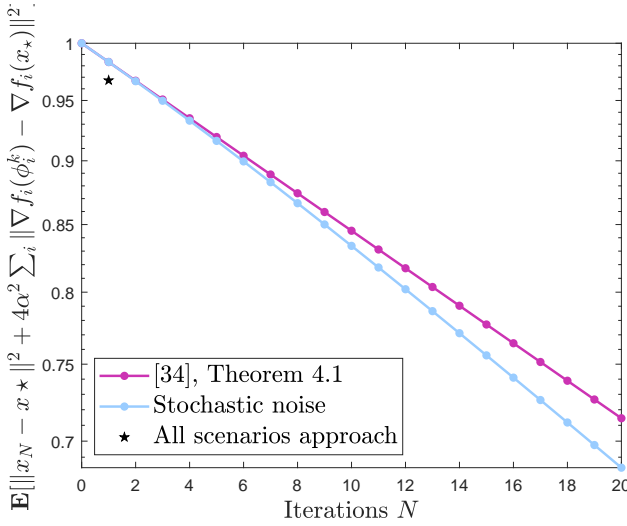


Fig. 8.5 Comparison of convergence analyses of (SGD) with constant stepsizes $\alpha = \frac{1}{6L}$, on $\mathcal{F}_{\mu,L}$, with noise satisfying $\mathbb{E}[\varepsilon_k | \mathcal{A}_k] = 0$, and (8.4) with parameters $A_1 = 4L$, $E_1 = 2$, $\rho = \frac{1}{n}$, and $A_2 = \frac{2L}{n}$. These assumptions hold for SAGA [DBLJ14]. We consider the performance measure $\mathbb{E}[\|x_N - x_\star\|^2 + 4n\alpha^2\sigma_N^2]$ with initial condition $\|x_0 - x_\star\|^2 + 4n\alpha^2\sigma_0^2 \leq 1$, for $\mu = 0.1$, $L = 1$, and $n = 10$. Displayed are the bound from [GHR20, Theorem 4.1] (pink), the PEP-based bound from $(P_{\text{relax},1})$ (blue), and the exact all-scenarios PEP-based bound (\star).

8.5 Concluding observations

We have proposed a framework for automated analysis of stochastic first-order methods, under all noise models satisfying Assumption 8.1, and independently of the noise's distribution. The framework positions itself between PEP analyses tailored for deterministic perturbation [THG17a, BTB22, DKGT20, VG25b, BTB22] (or equivalently, analyzing the worst realization of a stochastic perturbation) and those, regarding finite support settings, explicitly enumerating all scenarios [CKPG25, TB19, KHG23]. Indeed, by setting expected values as variables to PEP formulations, as in [AdKZ23b], our method improves upon deterministic analyses while maintaining similar complexity via relaxations such as $(P_{\text{relax},1})$. On the other hand, we provide an SDP formulation achieving exact worst-case bounds,

valid for all distributions satisfying the noise model, which is, in its implementation, similar to the all-scenarios approach, with 2^N scenarios of a specific equiprobable 2-realizations distribution. However, while the all-scenario approach yields convergence rates that are only valid for the setting considered, i.e., for the distribution analyzed, this exact formulation provides rates valid for the worst distribution possible with the same expectation and variance as the 2-realizations distribution, and thus for all such distributions, whose support may be infinite. This enables a unified framework applicable across diverse stochastic settings.

On structural noise models, e.g., finite-sum or block-coordinate settings, the proposed framework may yield looser rates than those obtained via the all-scenario approach, relying on this structural information. This is especially the case for problems with linear convergence rates. However, on problems with sublinear rates, it occasionally improves on the all-scenarios approach. Indeed, the latter is (mainly) restricted to analyses using potential functions, while our framework allows trading off tightness for complexity. That is, one can solve relaxations as $(P_{\text{relax},1})$, tractable on a large number of iterations, and that yield stronger convergence rates than those obtained via a potential function approach. Notably, the simplest relaxation often suffices to attain exact bounds.

Finally, the framework led to several noteworthy observations. Surprisingly, worst-case distributions within classes defined by expectation and variance are realized by simple, symmetric two-realizations distributions, orthogonal to all previous variables, which induce identical method trajectories. On the other hand, no insight into the structure of the associated worst-case function has been obtained. Additionally, as in the deterministic case, we showed that the worst-case analysis of stochastic methods reduces to a maximization problem over a convex set of size linear in N , the number of iterations, and independent of the algorithm's structure. However, we could not obtain any tractable representation for this convex set.

Appendix to Chapter 8

8.A Practical implementation of $(P_{\text{relax},1})$ and $(P_{\text{relax, sym}})$

This section provides the full SDP formulation for solving slightly generalized versions of (8.16) and (8.27). That is, we provide an example of SDP formulation for solving the simplest SDP relaxation $(P_{\text{relax},1})$ of (8.9), and its refinement $(P_{\text{relax, sym}})$, on the analysis of N iterations of (SGD) on $\mathcal{F}_{0,L}$, under the additive bounded noise model.

Consider first the general case of analyzing N iterations of a method \mathcal{M} on a class \mathcal{F} , under noise Ω_N , and given an initial condition C_{init} satisfying Assumption 8.1. In Section 8.3, we defined the Gram representation of sets $S = \{(x_i, f_i, g_i)\}_{i=0,1,\dots,N,\star} \in (\mathbb{R}^d \times \mathbb{R} \times \mathbb{R}^d)^{N+2}$, and $V = \{(\varepsilon_i, \sigma_i^2)\}_{i=0,1,\dots,N-1} \in (\mathbb{R}^d \times \mathbb{R})^N$ as

$$\begin{aligned} F &:= [f_\star, f_0, \dots, f_N, \sigma_0^2, \dots, \sigma_{N-1}^2]^\top, & (8.40) \\ G &:= PP^\top, \text{ where } P := [x_\star, x_0, \dots, x_N, g_\star, g_0, \dots, g_N, \varepsilon_0, \dots, \varepsilon_{N-1}]^\top, \end{aligned}$$

see (8.5). For computational purposes, we introduce a slightly different Gram representation of $S \cup V$, that is,

$$\begin{aligned} F &:= [f_0, \dots, f_N, \sigma_0^2, \dots, \sigma_{N-1}^2]^\top, & (8.41) \\ G &:= PP^\top, \text{ where } P := [x_0, g_0, \dots, g_N, \varepsilon_0, \dots, \varepsilon_{N-1}]^\top, \end{aligned}$$

This representation allows (i) reducing the size of the variables in all SDP formulations of (8.9), and (ii) encoding the optimal condition $\|g_\star\|^2 = 0$, and the method's dynamics directly via *selection vectors*, whose purpose are to access all scalar products in G , including those involving x_k , $k \geq 1$.

Specifically, we consider w.l.o.g. x_* to be the origin, and $f_* = 0$. To access elements of G , given \mathbf{e}_i the i^{th} -unit vector, we then define selection vectors

- $\mathbf{g}_* = \mathbf{x}_* = \mathbf{0} \in \mathbb{R}^{2N+2}$, $\mathbf{f}_* = \mathbf{0} \in \mathbb{R}^{2N+1}$,
- For $k = 0, 1, \dots, N$, $\mathbf{g}_k = \mathbf{e}_{k+2} \in \mathbb{R}^{2N+2}$ and $\mathbf{f}_k = \mathbf{e}_{k+1} \in \mathbb{R}^{2N+1}$,
- For $k = 0, 1, \dots, N-1$, $\mathbf{v}_k = \mathbf{e}_{N+3+k} \in \mathbb{R}^{2N+2}$ and $\mathbf{s}_k = \mathbf{e}_{N+2+k} \in \mathbb{R}^{2N+1}$,
- $\mathbf{x}_0 = \mathbf{e}_1 \in \mathbb{R}^{2N+2}$, and $\mathbf{x}_{k+1} = \mathbf{x}_0 + \sum_{i=0}^{k+1} \alpha_{k,i} (\mathbf{g}_i + \mathbf{v}_i)$, $k = 0, 1, \dots, N-1$.

These vectors allow accessing all f_i 's and σ_j^2 's ($i = 0, 1, \dots, N, \star, j = 0, 1, \dots, N-1$) in F , and scalar products of x_i, g_i, ε_j 's ($i = 0, 1, \dots, N, \star, j = 0, 1, \dots, N-1$) in G . For instance,

$$\begin{aligned} \mathbf{x}_i^\top G \mathbf{x}_j &= \langle x_i, x_j \rangle, & \mathbf{g}_i^\top G \mathbf{g}_j &= \langle g_i, g_j \rangle, & \mathbf{v}_i^\top G \mathbf{v}_j &= \langle \varepsilon_i, \varepsilon_j \rangle, \\ \mathbf{x}_i^\top G \mathbf{g}_j &= \langle x_i, g_j \rangle, & \mathbf{g}_i^\top G \mathbf{v}_j &= \langle g_i, \varepsilon_j \rangle, \\ \mathbf{x}_i^\top G \mathbf{v}_j &= \langle x_i, \varepsilon_j \rangle, \end{aligned}$$

and

$$F^\top (\mathbf{f}_i - \mathbf{f}_*) = f_i - f_*.$$

Using these selection vectors allows formulating (8.16) (generalized to N iterations) as follows.

$$\max_{(F,G) \in \mathbb{R}^{N+1} \times \mathcal{S}_+^{3N+2}} F^\top (\mathbf{f}_N - \mathbf{f}_*) \text{ s.t.}$$

$$\mathcal{F}\text{-interpolability: } F^\top (\mathbf{f}_i - \mathbf{f}_j) \geq \mathbf{g}_j^\top G (\mathbf{x}_i - \mathbf{x}_j) + \frac{1}{2L} (\mathbf{g}_j - \mathbf{g}_i)^\top G (\mathbf{g}_j - \mathbf{g}_i),$$

$$\forall i, j = 0, 1, \dots, N, \star,$$

$$C_{\text{init}}: (\mathbf{x}_0 - \mathbf{x}_*)^\top G (\mathbf{x}_0 - \mathbf{x}_*) \leq 1,$$

$$C_{\perp}: \mathbf{v}_k^\top G \mathbf{x}_j = \mathbf{v}_k^\top G \mathbf{g}_j = \mathbf{v}_k^\top G \mathbf{v}_l = 0, \quad k, l = 0, 1, \dots, N-1, \\ l \neq k, j \leq k,$$

$$C_{\text{variance}}: \mathbf{v}_k^\top G \mathbf{v}_k = \sigma^2, \quad k = 0, 1, \dots, N-1,$$

where satisfaction of \mathcal{M} arises from the definition of the selection vectors, and the interpolation conditions follow from (8.8).

Before providing the full formulation of (8.27), generalized to N iterations, we show that the constraint $\mathbf{G} \succcurlyeq 0$, in $(P_{\text{relax},\text{multi}})$, can be equivalently reduced to linear inequalities on G, G^{1j} .

Proposition 8.5. *Given $p, k \in \mathbb{N}$, $p \geq 1$, let $M \in (\mathbb{S}^k)^{2^p \times 2^p}$ be a 2^p -dimensional order p super-symmetric matrix whose elements belong to \mathbb{S}^k . Then, $M \succcurlyeq 0$ can be expressed via 3^p linear inequalities in the elements of M .*

Proof. We proceed by induction.

- If $p \geq 1$, $M \succcurlyeq 0$ if and only if $M_{1,1} \succcurlyeq 0$, $M_{1,1} - M_{1,2} \succcurlyeq 0$, $M_{1,1} + M_{1,2} \succcurlyeq 0$. Indeed,

$$\begin{aligned} \begin{pmatrix} M_{1,1} & M_{1,2} \\ M_{1,2} & M_{1,1} \end{pmatrix} \succcurlyeq 0 &\Leftrightarrow [x^\top \ y^\top] \begin{pmatrix} M_{1,1} & M_{1,2} \\ M_{1,2} & M_{1,1} \end{pmatrix} \begin{bmatrix} x \\ y \end{bmatrix} \succcurlyeq 0, \forall x, y \in \mathbb{R}^k \\ &\Leftrightarrow x^\top M_{1,1} x + y^\top M_{1,1} y + 2x^\top M_{1,2} y \geq 0, \forall x, y \in \mathbb{R}^k \\ &\Leftrightarrow (x+y)^\top (M_{1,1} + M_{1,2})(x+y) + (x-y)^\top (M_{1,1} - M_{1,2})(x-y) \geq 0, \\ &\quad \forall x, y \in \mathbb{R}^k. \end{aligned}$$

Condition $M_{1,1} \succcurlyeq 0$ follows from the case $y = 0$, $M_{1,1} - M_{1,2} \succcurlyeq 0$ from the case $y = x$, and $M_{1,1} + M_{1,2} \succcurlyeq 0$ from the case $y = -x$.

- Suppose the assertion is true for p , and consider the case $p + 1$. Partition M into four blocks in $(\mathbb{S}^k)^{2^p \times 2^p}$:

$$M = \begin{pmatrix} \bar{M}_{1,1} & \bar{M}_{1,2} \\ \bar{M}_{1,2} & \bar{M}_{1,1} \end{pmatrix}.$$

Then, $M \succcurlyeq 0$ if and only if $\bar{M}_{1,1} \succcurlyeq 0$, $\bar{M}_{1,1} - \bar{M}_{1,2} \succcurlyeq 0$, $\bar{M}_{1,1} + \bar{M}_{1,2} \succcurlyeq 0$. By assumption, each of these 3 positive definiteness constraints can be expressed via 3^p linear expressions in the elements (in \mathbb{R}^k) of $\bar{M}_{1,1}$, $\bar{M}_{1,1} - \bar{M}_{1,2}$, and $\bar{M}_{1,1} + \bar{M}_{1,2}$, respectively, as all are order p super-symmetric. The total number of constraints is thus 3^{p+1} .

□

Building on Proposition 8.5, it holds that the N iterations version of (8.22) can be expressed as

$$\begin{aligned} &\max_{\substack{(F,G) \in \mathbb{R}^{N+1} \times \mathbb{S}^{2N+2}, \\ G^{12} \in \mathbb{S}^{2N+2}}} F^\top (\mathbf{f}_N - \mathbf{f}_\star) \text{ s.t.} \\ \begin{pmatrix} G & G^{12} \\ G^{12} & G \end{pmatrix} \succcurlyeq 0: & \quad G \succcurlyeq 0, G + G^{12} \succcurlyeq 0, G - G^{12} \succcurlyeq 0, \\ \mathcal{F}\text{-interpolability:} & \quad F^\top (\mathbf{f}_i - \mathbf{f}_j) \geq \mathbf{g}_j^\top G (\mathbf{x}_i - \mathbf{x}_j) \end{aligned}$$

$$\begin{aligned}
 & + \frac{1}{2L} (\mathbf{g}_j - \mathbf{g}_i)^\top G (\mathbf{g}_j - \mathbf{g}_i), \quad i, j = 0, 1, \dots, N, \star, \\
 & F^\top (\mathbf{f}_i - \mathbf{f}_j) \geq \mathbf{g}_j^\top G^{12} \mathbf{x}_i - \mathbf{g}_j^\top G \mathbf{x}_j \\
 & + \frac{1}{2L} \left(\mathbf{g}_j^\top G \mathbf{g}_j + \mathbf{g}_i^\top G \mathbf{g}_i - 2\mathbf{g}_j^\top G^{12} \mathbf{g}_j \right), \\
 & i, j = 0, 1, \dots, N, \star, \\
 C_{\text{init}}: & \quad (\mathbf{x}_0 - \mathbf{x}_\star)^\top G (\mathbf{x}_0 - \mathbf{x}_\star) \leq 1, \\
 C_{\perp}: & \quad \mathbf{v}_k^\top G \mathbf{x}_j = \mathbf{v}_k^\top G \mathbf{g}_j = \mathbf{v}_k^\top G \mathbf{v}_l = 0, \quad \forall l \neq k, j \leq k, \\
 C_{\text{variance}}: & \quad \mathbf{v}_k^\top G \mathbf{v}_k = \sigma^2, \quad k = 0, 1, \dots, N-1, \\
 \text{Covariance:} & \quad \mathbf{v}_k^\top G^{12} \mathbf{v}_k = -\sigma^2, \quad k = 0, 1, \dots, N-1,
 \end{aligned}$$

where semi-definite positiveness constraint follows from Proposition 8.5.

PART IV
Conclusion

9

Research outcomes and perspectives

This thesis proposed new tools to tackle open challenges in deriving interpolation conditions, that is, exact discrete characterizations of function and operator classes. In particular, we established such characterizations (or, at least, refined existing ones) for a few classes that had previously lacked tight discrete descriptions. We also investigated the question of exactly describing expectations with respect to random variables, in classes of distributions with specified variance and expectation. These developments enabled several extensions of computer-aided exact performance analysis. In particular, we proposed extensions of the Performance Estimation Problem (PEP) framework to weakly convex optimization, second-order univariate optimization, and stochastic first-order optimization, including methods operating under noise models with infinite support.

We now summarize the main research outcomes of this research, and overview some perspectives and open challenges.

9.1 Research outcomes

Advances in function and operator interpolation Chapters 4 and 5 proposed a constraint-based, algebraic approach to derive interpolation conditions for function and operator classes in a principled manner.

More specifically, Chapter 4 proposed treating constraints themselves as primary objects in interpolation theory, defining function or operator classes algebraically by imposing these constraints globally. This allows highlighting classes based on whether they admit a simple discrete description rather than on their analytic properties.

Additionally, Chapter 4 introduced the notion of pointwise extensibility (Definition 4.5), providing an algebraic framework to assess the interpolability of a constraint—whether derived from a function class or within the constraint-based approach—by (i) automatically generating counterexamples to interpolability, or (ii) reducing interpolability proofs to algebraic questions independent of the class’s analytic properties.

Further, Chapter 5 introduced a constructive procedure to derive interpolation conditions for a given class. Starting from any description of a function or operator class, it iteratively refines the characterization, producing interpolation conditions once stable. This approach strengthens sets of conditions based only on algebraic properties, yielding improved conditions that would have been challenging to guess based solely on the class’ analytical properties. Proving a constraint’s interpolability reduces to showing its stability under this procedure, i.e., it can no longer be strengthened.

Finally, Chapter 6 focused on interpolation for *univariate* function classes. It introduced a generic, principled approach to lift interpolation conditions from a function class \mathcal{F} to the class of functions whose derivatives lie in \mathcal{F} . This method enables the derivation of interpolation conditions for a wide range of classes, particularly those defined by higher-order properties. Although inherently limited to univariate functions, the resulting conditions offer valuable insights into the structure of multivariate interpolation conditions.

We briefly summarize several answers to the open questions raised in Chapter 2, as proposed throughout this work.

Summary

Question 1. Given a function/operator class \mathcal{F} , how to guess a candidate interpolation condition for \mathcal{F} ?

- Starting from a given discrete characterization of \mathcal{F} , rely on the constructive strengthening procedure introduced in Chapter 5.
- Using the approach presented in Chapter 1, obtain interpolation conditions for the univariate restriction of \mathcal{F} . Building on these interpolation conditions, guess a multivariate candidate.

Question 2. How to establish a constraint's interpolability for \mathcal{F} ?

- Use pointwise extensibility as introduced in Chapter 4 to establish the interpolability of the constraint algebraically.
- Show stability of the constraint under the strengthening procedure proposed in Chapter 5.

And to answer both questions at once?

- Use the constraint-based approach proposed in Chapter 4 and obtain, altogether, function/operator classes and their interpolation conditions.

Interpolation (or improved) conditions for a few first-order problem classes

Chapter 4 applied pointwise extensibility to derive interpolation conditions for the class of weakly convex functions with bounded subgradients. These conditions were shown to be Gram-representable, enabling a tight numerical analysis of the subgradient method on this class.

Chapter 5 then refined existing descriptions of several problem classes, including those of smooth functions satisfying a Łojasiewicz condition, blockwise smooth convex functions, strongly monotone Lipschitz operators, and uniformly convex functions. The strengthened conditions preserve the structure of the originals: Gram-representable conditions remain SDP-representable. These improved descriptions enabled sharper numerical convergence bounds when used within PEP formulations.

Interpolation conditions for second-order univariate optimization Building on Chapter 6, Chapter 7 took a first step towards exact analysis of

second-order methods by extending the PEP framework to the tight analysis of *univariate* optimization with second-order information. This was achieved by (i) deriving exact discrete characterizations of several univariate function classes defined via second-order properties, and (ii) proposing tractable second-order PEP formulations using recent progress on non-convex PEPs.

These developments provided insight into the structure of interpolation conditions for second-order function classes, which are already non-trivial in the univariate case. They also enabled the derivation of exact convergence rates in this setting, which improve upon existing univariate bounds, and serve as lower bounds in the multivariate case. Moreover, when these match known multivariate rates, they certify their tightness by exhibiting explicit worst-case instances. Finally, the computer-aided framework facilitates fair numerical comparisons across methods.

Interpolation conditions for stochastic first-order optimization Finally, Chapter 8 introduced a framework for the automated analysis of stochastic first-order methods. Given a method, a function class, and a noise model specified by its expectation and variance, the framework formulates a hierarchy of SDP problems that yield increasingly tight convergence guarantees. These are based on a range of necessary, and ultimately sufficient, conditions characterizing the expected values, under any distribution in the noise model, arising in the method. These expected values are embedded as variables in the PEP formulations.

The simplest such formulation, whose complexity scales linearly with the number of iterations, provides convergence rates that turn out, numerically, to be already exact on many problem classes. This tightness is verified, for a small number of iterations, by solving the exact PEP formulation, whose complexity scales exponentially with the number of iterations.

The framework accommodates various stochastic settings, with finite or infinite support, including the unstructured noise model with bounded variance, finite-sum optimization, and block-coordinate methods, in a unified manner, as guarantees apply to any setting consistent with the noise model, i.e., its expectation and variance. This unified approach comes at the cost of yielding rates which, though exact under the noise model, may be suboptimal on specific noise structures, e.g., finite-sum optimization or block-coordinate methods. Using this structural information via, e.g., an all-scenarios approach [TB19], yields, in theory, better convergence rates. However, since the all-scenarios approach scales exponentially with

the number of iterations and is hence restricted to the analysis of a small number of iterations, the scalable small-sized relaxation of our framework sometimes improves on all-scenarios convergence guarantees.

9.2 Perspectives and open questions

The tools introduced in this thesis open several avenues for future research. On the one hand, they can be further leveraged to extend the Performance Estimation Problem (PEP) framework to broader settings. On the other hand, they also come with certain limitations, in particular, their practical tractability. These limitations raise open challenges that remain to be addressed.

Further exploitation of the proposed tools

The tools proposed throughout this thesis can be readily extended to analyze a wide range of settings beyond those considered in this work. For instance, the refinement procedure introduced in Chapter 5 can, in principle, be applied to strengthen any non-sufficient constraint.

Exploitation of the univariate framework of Chapters 6 and 7 The univariate interpolation framework developed in Chapter 6 can be leveraged to derive interpolation conditions for additional function classes, including those characterized by higher-order properties. These conditions can then be used within the PEP framework to tightly analyze a broad range of algorithms in the univariate setting, such as tensor methods [Nes21], zeroth-order methods, or adaptive methods. Such analyses would provide insight into the bottlenecks of the method, exact bounds in the univariate setting, and lower bounds in the multivariate setting.

Exploitation of the stochastic framework of Chapter 8 Similarly, while the stochastic framework of Chapter 8 was primarily used to analyze (SGD) on the class $\mathcal{F}_{\mu,L}$ of smooth (strongly) convex functions, the approach naturally extends to a wider array of methods and function classes. As discussed in Section 8.3, any problem satisfying Assumption 8.1, including accelerated methods, and those applied to non-convex or non-smooth problems, can be addressed using this framework. Beyond analyzing existing

methods, the framework may also be used for the principled design of new algorithms, as illustrated in [DT20].

Finally, the proposed framework supports a reverse perspective: identifying the noise assumptions under which desirable convergence rates are achievable, and designing algorithms that satisfy them.

Analytical proofs Finally, while the thesis primarily relied on solving, numerically, PEP formulations under their primal form, one could also consider dual formulations, either to obtain dual numerical results or to retrieve analytical guarantees, as introduced in Chapter 3.

Open questions

On the limits of the algebraic approach to interpolation The algebraic approach to interpolation developed in Chapters 4 and 5 presents several practical limitations. In particular, the resulting proofs, using this algebraic framework, are often tedious and require familiarity with non-trivial algebraic concepts. This drawback can be partially mitigated by leveraging symbolic computation tools, such as *Mathematica* notebooks [Inc].

Further, the one-point strengthening procedure presented in Chapter 5, which allows, theoretically, obtaining interpolation conditions for any given function or operator class in an automated way, is in practice limited. For instance, one-point strengthenings are expressed as minimax optimization problems, whose resolution is in general challenging. In particular, the tools we used for solving such problems, which consist of dualizing the inner problem, fail on non-convex initial constraints, and may only provide relaxed strengthenings. Moreover, due to the complexity of solving this minimax problem, we are quickly limited in the iterative procedure, since formulations become hardly manipulable, both for computing their one-point strengthening and for proving their stability. Hence, in practice, (i) starting from wisely defined constraints is, in fact, crucial to obtain interpolation conditions as final products, and (ii) proving stability of a constraint, hence its interpolability, is almost never feasible.

The many remaining open challenges thus include the question of tackling conditions whose one-point strengthenings are not convex, by relying on tools other than dualization. Another venue for improvement is that of specializing the approach to more specific classes of problems. Yet another possible improvement could be to search for one-point strengthenings without relaxations, for instance, by relying on quantifier elimination

techniques (see, e.g., [BPR06]), that do not scale well in the number of variables and parameters in general, but that could handle our problems with very few of them.

Structure of interpolation conditions In Chapter 5, we observed that some interpolation conditions are not pairwise, i.e., they involve more than two points. In particular, we obtained refined descriptions of some function and operator classes, involving three points. However, we were not able to conclude on the interpolability of these refinements. This leaves the following questions open: are there function/operator classes whose interpolation conditions are *terwise*? Or involving even more datapoints? In addition, guessing in advance which classes admit pairwise interpolation conditions is also an open challenge.

Computer-aided analysis of multivariate second-order optimization The extension of the PEP framework to the analysis of second-order methods presented in Chapter 7 is limited to univariate analyses. A natural extension would thus be to tackle the multivariate case. This would require (i) obtaining multivariate interpolation conditions, for instance by building on the interpolation conditions we derived, which naturally furnish candidate multivariate conditions-but guessing the right candidate is a priori not straightforward, and (ii) finding a way to efficiently solve PEP formulations involving matrix variables accounting for the Hessians, an already computationally expensive task in the univariate setting.

Limitations of the stochastic framework The main limitation of the stochastic framework proposed in Chapter 8 is its scalability to obtain tight convergence guarantees, since exact rates only result from formulations whose size grows exponentially with N , the number of iterations analyzed.

An important avenue for improvement lies in reducing the computational complexity of the framework. One direction would be to derive an exact description of the convex set \mathcal{G}_E of pairs (F, G) corresponding to actual problems, with size scaling linearly in N , or to further rely on the symmetry of the exponential formulation. Alternatively, building on the observation that the simplest relaxation in the proposed framework is already tight for many problem classes, one could develop an automated procedure to *verify* whether this relaxation yields exact guarantees, or better understand what problems possess this tightness property.

Are we analyzing the right problem classes? Finally, a more prospective question that arises from all observations made in this work is the following: are the function and operator classes classically used in optimization theory the right ones? Indeed, one would expect function and operator classes of interest to have simple discrete characterizations, since any performance analysis builds on those discrete characterizations rather than global ones. However, it seems that a range of classical function classes, especially second-order function classes, admit complex discrete characterizations.

A less immediate, but promising, direction would thus be, for instance, to further exploit the constraint-based approach proposed in Chapter 4. In particular, analyzing a broader set of candidate constraints to identify those that are interpolable and define meaningful function or operator classes could lead to the discovery of new classes. These classes would be characterized by the simplicity of their interpolation conditions, an especially appealing prospect in the context of second-order analysis.

Bibliography

- [Abb24] Hadi Abbaszadehpeivasti. *Performance analysis of optimization methods for machine learning*. PhD thesis, Tilburg University, 2024.
- [AD19] Hilal Asi and John C Duchi. Stochastic (approximate) proximal point methods: Convergence, optimality, and adaptivity. *SIAM Journal on Optimization*, 29(3):2257–2290, 2019.
- [AdKZ23a] Hadi Abbaszadehpeivasti, Etienne de Klerk, and Moslem Zamani. Conditions for linear convergence of the gradient method for non-convex optimization. *Optimization Letters*, 17(5):1105–1125, 2023.
- [AdKZ23b] Hadi Abbaszadehpeivasti, Etienne de Klerk, and Moslem Zamani. Convergence rate analysis of randomized and cyclic coordinate descent for convex optimization through semidefinite programming. *Applied Set-Valued Analysis and Optimization*, 5(2):141–153, 2023.
- [AdKZ24] 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, 2024.
- [AM17a] Daniel Azagra and Carlos Mudarra. An extension theorem for convex functions of class $C^{1,1}$ on hilbert spaces. *Journal of Mathematical Analysis and Applications*, 446(2):1167–1182, 2017.
- [AM17b] Daniel Azagra and Carlos Mudarra. Whitney extension theorems for convex functions of the classes C^1 and $C^{1,\omega}$. *Proceedings of the London Mathematical Society*, 114(1):133–158, 2017.

- [AM19] Daniel Azagra and Carlos Mudarra. Smooth convex extensions of convex functions. *Calculus of Variations and Partial Differential Equations*, 58(84), 2019.
- [AP24] Jason M Altschuler and Pablo A Parrilo. Acceleration by step-size hedging: Silver stepsize schedule for smooth convex optimization. *Mathematical Programming*, pages 1–14, 2024.
- [AZ18] Zeyuan Allen-Zhu. Katyusha: The first direct acceleration of stochastic gradient methods. *Journal of Machine Learning Research*, 18(221):1–51, 2018.
- [Bac10] Francis Bach. Self-concordant analysis for logistic regression. *Electronic Journal of Statistics*, 4:384 – 414, 2010.
- [Bar21] Mathieu Barré. *Worst-case analysis of efficient first-order methods*. PhD thesis, Université Paris sciences et lettres, 2021.
- [BB07] Léon Bottou and Olivier Bousquet. The tradeoffs of large scale learning. *Advances in neural information processing systems*, 20, 2007.
- [BC17] Heinz H Bauschke and Patrick L Combettes. *Convex analysis and monotone operator theory in Hilbert spaces*. Springer, 2017.
- [BCN18] Léon Bottou, Frank E Curtis, and Jorge Nocedal. Optimization methods for large-scale machine learning. *SIAM review*, 60(2):223–311, 2018.
- [BDL07] Jérôme Bolte, Aris Daniilidis, and Adrian Lewis. The Łojasiewicz inequality for nonsmooth subanalytic functions with applications to subgradient dynamical systems. *SIAM Journal on Optimization*, 17(4):1205–1223, 2007.
- [Bec17] Amir Beck. *First-order methods in optimization*. SIAM, 2017.
- [Ber09] Dimitri Bertsekas. *Convex optimization theory*, volume 1. Athena Scientific, 2009.
- [BHG24] Nizar Bousselmi, Julien M Hendrickx, and François Glineur. Interpolation conditions for linear operators and applications to performance estimation problems. *SIAM Journal on Optimization*, 34(3):3033–3063, 2024.

- [BNPS17] Jérôme Bolte, Trong Phong Nguyen, Juan Peypouquet, and Bruce W Suter. From error bounds to the complexity of first-order descent methods for convex functions. *Mathematical Programming*, 165:471–507, 2017.
- [Bor87] Karl-Heinz Borgwardt. The simplex algorithm: a probabilistic analysis. *Algorithms and Combinatorics*, 1, 1987.
- [Bor08] Vivek S Borkar. *Stochastic approximation: a dynamical systems viewpoint*, volume 9. Springer, 2008.
- [Boy04] Stephen Boyd. Convex optimization. *Cambridge UP*, 2004.
- [BPR06] Saugata Basu, Richard M. Pollack, and Marie-Françoise Roy. *Algorithms in Real Algebraic Geometry*, volume 10 of *Algorithms and Computation in Mathematics*. Springer-Verlag, Berlin, 2006.
- [BS00] Robert G Bartle and Donald R Sherbert. *Introduction to real analysis*. John Wiley & Sons, Inc., New Jersey, 2000.
- [BS01] Yuri Brudnyi and Pavel Shvartsman. Whitney’s extension problem for multivariate $c^{1,w}$ -functions. *Transactions of the American Mathematical Society*, 353(6):2487–2512, 2001.
- [BST14] Raef Bassily, Adam Smith, and Abhradeep Thakurta. Private empirical risk minimization: Efficient algorithms and tight error bounds. In *2014 IEEE 55th annual symposium on foundations of computer science*, pages 464–473. IEEE, IEEE, 2014.
- [BT08] Dimitri Bertsekas and John N Tsitsiklis. *Introduction to probability*, volume 1. Athena Scientific, 2008.
- [BT13] Amir Beck and Luba Tetruashvili. On the convergence of block coordinate descent type methods. *SIAM Journal on Optimization*, 23(4):2037–2060, 2013.
- [BTB22] Mathieu Barré, Adrien B Taylor, and Francis Bach. Principled analyses and design of first-order methods with inexact proximal operators. *Mathematical Programming*, pages 1–46, 2022.
- [BTd20] Mathieu Barré, Adrien Taylor, and Alexandre d’Aspremont. Complexity guarantees for polyak steps with momentum. In *Conference on learning theory*, pages 452–478. PMLR, 2020.

- [BTN02] Aharon Ben-Tal and Arkadi Nemirovski. Robust optimization—methodology and applications. *Mathematical programming*, 92:453–480, 2002.
- [Bur85] James V Burke. Descent methods for composite nondifferentiable optimization problems. *Mathematical Programming*, 33:260–279, 1985.
- [BW21] Axel Böhm and Stephen J Wright. Variable smoothing for weakly convex composite functions. *Journal of optimization theory and applications*, 188:628–649, 2021.
- [C⁺47] Augustin Cauchy et al. Méthode générale pour la résolution des systèmes d'équations simultanées. *Comp. Rend. Sci. Paris*, 25(1847):536–538, 1847.
- [CDO18] Michael Cohen, Jelena Diakonikolas, and Lorenzo Orecchia. On acceleration with noise-corrupted gradients. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 1019–1028. PMLR, 10–15 Jul 2018.
- [CGT00] Andrew R Conn, Nicholas IM Gould, and Philippe L Toint. Trust region methods. *SIAM*, 2000.
- [CGT10] Coralia Cartis, Nicholas IM Gould, and Ph L 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.
- [CGT11] Coralia Cartis, Nicholas IM Gould, and Philippe L Toint. Adaptive cubic regularisation methods for unconstrained optimization. Part I: motivation, convergence and numerical results. *Mathematical Programming*, 127(2):245–295, 2011.
- [CH23] Sebastien Colla and Julien M Hendrickx. Automatic performance estimation for decentralized optimization. *IEEE Transactions on Automatic Control*, 68(12):7136–7150, 2023.

- [CHVSL18] Saman Cyrus, Bin Hu, Bryan Van Scoy, and Laurent Lessard. A robust accelerated optimization algorithm for strongly convex functions. In *2018 Annual American Control Conference (ACC)*, pages 1376–1381. IEEE, IEEE, 2018.
- [CKPG25] Daniel Cortild, Lucas Ketels, Juan Peypouquet, and Guillaume Garrigos. New tight bounds for sgd without variance assumption: A computer-aided lyapunov analysis. *arXiv preprint arXiv:2505.17965*, 2025.
- [Cla36] James A. Clarkson. Uniformly convex spaces. *Transactions of the American Mathematical Society*, 40(3):396–414, 1936.
- [Col24] Sébastien Colla. *Computer-Aided Analysis of Decentralized Optimization Methods*. PhD thesis, UCL-Université Catholique de Louvain, 2024.
- [d’A08] Alexandre d’Aspremont. Smooth optimization with approximate gradient. *SIAM Journal on Optimization*, 19(3):1171–1183, 2008.
- [DBLJ14] Aaron Defazio, Francis Bach, and Simon Lacoste-Julien. Saga: A fast incremental gradient method with support for non-strongly convex composite objectives. *Advances in neural information processing systems*, 27, 2014.
- [DD19] Damek Davis and Dmitriy Drusvyatskiy. Stochastic model-based minimization of weakly convex functions. *SIAM Journal on Optimization*, 29(1):207–239, 2019.
- [DDG⁺22] Marina Danilova, Pavel Dvurechensky, Alexander Gasnikov, Eduard Gorbunov, Sergey Guminov, Dmitry Kamzolov, and Innokentiy Shibaev. Recent theoretical advances in non-convex optimization. In *High-Dimensional Optimization and Probability: With a View Towards Data Science*, pages 79–163. Springer, 2022.
- [DDMP18] Damek Davis, Dmitriy Drusvyatskiy, Kellie J MacPhee, and Courtney Paquette. Subgradient methods for sharp weakly convex functions. *Journal of Optimization Theory and Applications*, 179:962–982, 2018.

- [DG19] Damek Davis and Benjamin Grimmer. Proximally guided stochastic subgradient method for nonsmooth, nonconvex problems. *SIAM Journal on Optimization*, 29(3):1908–1930, 2019.
- [DGFST24] 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*, 2024.
- [DGN13] Olivier Devolder, François Glineur, and Yurii Nesterov. First-order methods with inexact oracle: the strongly convex case. *CORE Discussion Papers*, 2013016:47, 2013.
- [DGN14] Olivier Devolder, François Glineur, and Yurii Nesterov. First-order methods of smooth convex optimization with inexact oracle. *Mathematical Programming*, 146:37–75, 2014.
- [DGVPR24] Shuvomoy Das Gupta, Bart PG Van Parys, and Ernest K Ryu. Branch-and-bound performance estimation programming: A unified methodology for constructing optimal optimization methods. *Mathematical Programming*, 204(1):567–639, 2024.
- [DHLGL18] Aris Daniilidis, Mounir Haddou, Erwan Le Gruyer, and Olivier Ley. Explicit formulas for $C^{1,1}$ Glaeser-Whitney extensions of 1-Taylor fields in Hilbert spaces. *Proceedings of the American Mathematical Society*, 146(10):4487–4495, 2018.
- [DKGT17] Etienne De Klerk, François Glineur, and Adrien B Taylor. On the worst-case complexity of the gradient method with exact line search for smooth strongly convex functions. *Optimization Letters*, 11:1185–1199, 2017.
- [DKGT20] Etienne De Klerk, Francois Glineur, and Adrien B Taylor. Worst-case convergence analysis of inexact gradient and newton methods through semidefinite programming performance estimation. *SIAM Journal on Optimization*, 30(3):2053–2082, 2020.
- [DMN24] Nikita Doikov, Konstantin Mishchenko, and Yurii Nesterov. Super-universal regularized Newton method. *SIAM Journal on Optimization*, 34(1):27–56, 2024.

- [DN21] Nikita Doikov and Yurii Nesterov. Minimizing uniformly convex functions by cubic regularization of newton method. *Journal of Optimization Theory and Applications*, 189(1):317–339, 2021.
- [DN24] Nikita Doikov and Yurii Nesterov. Gradient regularization of newton method with bregman distances. *Mathematical programming*, 204(1):1–25, 2024.
- [Doi21] Nikita Doikov. *New second-order and tensor methods in convex optimization*. PhD thesis, Université catholique de Louvain, Louvain-la-Neuve, Belgium, 2021.
- [Doi23] Nikita Doikov. Minimizing quasi-self-concordant functions by gradient regularization of Newton method. *arXiv preprint arXiv:2308.14742*, 2023.
- [DP19] Dmitriy Drusvyatskiy and Courtney Paquette. Efficiency of minimizing compositions of convex functions and smooth maps. *Mathematical Programming*, 178:503–558, 2019.
- [Dro17] Yoel Drori. The exact information-based complexity of smooth convex minimization. *Journal of Complexity*, 39:1–16, 2017.
- [Dro20] Yoel Drori. On the properties of convex functions over open sets. *Journal of Convex Analysis*, 27:1303–1314, 2020.
- [DS09] John Duchi and Yoram Singer. Efficient online and batch learning using forward backward splitting. *The Journal of Machine Learning Research*, 10:2899–2934, 2009.
- [DS20] Yoel Drori and Ohad Shamir. The complexity of finding stationary points with stochastic gradient descent. In *International Conference on Machine Learning*, pages 2658–2667. PMLR, 2020.
- [DT14] Yoel Drori and Marc Teboulle. Performance of first-order methods for smooth convex minimization: a novel approach. *Mathematical Programming*, 145(1):451–482, 2014.
- [DT16] Yoel Drori and Marc Teboulle. An optimal variant of kelley’s cutting-plane method. *Mathematical Programming*, 160(1):321–351, 2016.

- [DT20] Yoel Drori and Adrien B Taylor. Efficient first-order methods for convex minimization: a constructive approach. *Mathematical Programming*, 184(1-2):183–220, 2020.
- [DTdB22] Radu-Alexandru Dragomir, Adrien B Taylor, Alexandre d’Aspremont, and Jérôme Bolte. Optimal complexity and certification of Bregman first-order methods. *Mathematical Programming*, 194:41–83, 2022.
- [Dur19] Rick Durrett. *Probability: theory and examples*, volume 49. Cambridge university press, 2019.
- [Fef05] Charles L Fefferman. A sharp form of whitney’s extension theorem. *Annals of mathematics*, 161(1):509–577, 2005.
- [Fic64] Gaetano Fichera. *Problemi elastostatici con vincoli unilaterali: il problema di Signorini con ambigue condizioni al contorno*. Accademia nazionale dei Lincei, 1964.
- [FIL17] Charles Fefferman, Arie Israel, and Garving K Luli. Interpolation of data by smooth nonnegative functions. *Revista matemática iberoamericana*, 33(1):305–324, 2017.
- [FR15] Olivier Fercoq and Peter Richtárik. Accelerated, parallel, and proximal coordinate descent. *SIAM Journal on Optimization*, 25(4):1997–2023, 2015.
- [G⁺14] Ian Goodfellow et al. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, pages 2672–2680, 2014.
- [Gan22] Oran Gannot. A frequency-domain analysis of inexact gradient methods. *Mathematical Programming*, 194(1-2):975–1016, 2022.
- [GDT23] Baptiste Goujaud, Aymeric Dieuleveut, and Adrien Taylor. On fundamental proof structures in first-order optimization. In *2023 62nd IEEE Conference on Decision and Control (CDC)*, pages 3023–3030. IEEE, IEEE, 2023.
- [GEIGM22] Charles Guille-Escuret, Adam Ibrahim, Baptiste Goujaud, and Ioannis Mitliagkas. Gradient descent is optimal under lower restricted secant inequality and upper error bound. *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.

- [GG23] Guillaume Garrigos and Robert M Gower. Handbook of convergence theorems for (stochastic) gradient methods. *arXiv preprint arXiv:2301.11235*, 2023.
- [GHR20] Eduard Gorbunov, Filip Hanzely, and Peter Richtárik. A unified theory of sgd: Variance reduction, sampling, quantization and coordinate descent. In *International Conference on Artificial Intelligence and Statistics*, pages 680–690. PMLR, 2020.
- [GL12] Saeed Ghadimi and Guanghui Lan. Optimal stochastic approximation algorithms for strongly convex stochastic composite optimization i: A generic algorithmic framework. *SIAM Journal on Optimization*, 22(4):1469–1492, 2012.
- [GL13] Saeed Ghadimi and Guanghui Lan. Stochastic first-and zeroth-order methods for nonconvex stochastic programming. *SIAM journal on optimization*, 23(4):2341–2368, 2013.
- [Gla58] Georges Glaeser. Étude de quelques algèbres tayloriennes. *Journal d’analyse mathématique*, 6(1):1–124, 1958.
- [GLG22] Eduard Gorbunov, Nicolas Loizou, and Gauthier Gidel. Extragradient method: $O(1/k)$ last-iterate convergence for monotone variational inequalities and connections with cocoercivity. In *International Conference on Artificial Intelligence and Statistics*, pages 366–402. PMLR, 2022.
- [GLQ⁺19] Robert Mansel Gower, Nicolas Loizou, Xun Qian, Alibek Sailanbayev, Egor Shulgin, and Peter Richtárik. Sgd: General analysis and improved rates. In *International conference on machine learning*, pages 5200–5209. PMLR, 2019.
- [GMG⁺22] Baptiste Goujaud, Céline Mouter, François Glineur, Julien Hendrickx, Adrien Taylor, and Aymeric Dieuleveut. PEPit: computer-assisted worst-case analyses of first-order optimization methods in Python. *arXiv preprint arXiv:2201.04040*, 2022.
- [GMG⁺24] Baptiste Goujaud, Céline Mouter, François Glineur, Julien M. Hendrickx, Adrien B. Taylor, and Aymeric Dieuleveut. Pepit: computer-assisted worst-case analyses of first-order optimization methods in Python. *Mathematical Programming Computation*, 16(3):337–367, 2024.

- [Gri24] Benjamin Grimmer. Provably faster gradient descent via long steps. *SIAM Journal on Optimization*, 34(3):2588–2608, 2024.
- [Gru09] Erwan Le Gruyer. Minimal Lipschitz extensions to differentiable functions defined on a Hilbert space. *Geometric and Functional Analysis*, 19(4):1101–1118, 2009.
- [GSBR20] Robert M Gower, Mark Schmidt, Francis Bach, and Peter Richtárik. Variance-reduced methods for machine learning. *Proceedings of the IEEE*, 108(11):1968–1983, 2020.
- [GTD22] Baptiste Goujaud, Adrien Taylor, and Aymeric Dieuleveut. Optimal first-order methods for convex functions with a quadratic upper bound. *preprint arXiv:2205.15033*, 2022.
- [GTD23] Baptiste Goujaud, Adrien Taylor, and Aymeric Dieuleveut. Provable non-accelerations of the heavy-ball method. *arXiv preprint arXiv:2307.11291*, 2023.
- [GTG22] Eduard Gorbunov, Adrien Taylor, and Gauthier Gidel. Last-iterate convergence of optimistic gradient method for monotone variational inequalities. *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- [GTHG23] Eduard Gorbunov, Adrien Taylor, Samuel Horváth, and Gauthier Gidel. Convergence of proximal point and extragradient-based methods beyond monotonicity: the case of negative comonotonicity. In *International Conference on Machine Learning*, pages 11614–11641. PMLR, 2023.
- [Gur23] Gurobi Optimization, LLC. Gurobi Optimizer Reference Manual, 2023.
- [GY20] Guoyong Gu and Junfeng Yang. Tight sublinear convergence rate of the proximal point algorithm for maximal monotone inclusion problems. *SIAM Journal on Optimization*, 30(3):1905–1921, 2020.
- [GY24] Guoyong Gu and Junfeng Yang. Tight ergodic sublinear convergence rate of the relaxed proximal point algorithm for monotone variational inequalities. *Journal of Optimization Theory and Applications*, 202(1):373–387, 2024.

- [Hil88] David Hilbert. Über die darstellung definiter formen als summe von formenquadraten. *Mathematische Annalen*, 32(3):342–350, 1888.
- [Hil21] Roland Hildebrand. Optimal step length for the Newton method: case of self-concordant functions. *Mathematical Methods of Operations Research*, 94(2):253–279, 2021.
- [HKP⁺22] Slavomír Hanzely, Dmitry Kamzolov, Dmitry Pasechnyuk, Alexander Gasnikov, Peter Richtárik, and Martin Takáč. A damped Newton method achieves global $\mathcal{O}(1/k^2)$ and local quadratic convergence rate. *Advances in Neural Information Processing Systems*, 35:25320–25334, 2022.
- [HL17] Bin Hu and Laurent Lessard. Dissipativity theory for nesterov’s accelerated method. In *International Conference on Machine Learning*, pages 1549–1557. PMLR, 2017.
- [Höl86] Otto Hölder. Ueber die eigenschaft der gammafunction keiner algebraischen differentialgleichung zu genügen. *Mathematische Annalen*, 28(1):1–13, 1886.
- [HSL21] Bin Hu, Peter Seiler, and Laurent Lessard. Analysis of biased stochastic gradient descent using sequential semidefinite programs. *Mathematical programming*, 187:383–408, 2021.
- [HSR17] Bin Hu, Peter Seiler, and Anders Rantzer. A unified analysis of stochastic optimization methods using jump system theory and quadratic constraints. In *Conference on learning theory*, pages 1157–1189. PMLR, 2017.
- [IH24] Anastasia Ivanova and Roland Hildebrand. Optimal step length for the maximal decrease of a self-concordant function by the Newton method. *Optimization Letters*, 18(3):847–854, 2024.
- [Inc] Wolfram Research, Inc. Mathematica, Version 14.2. Champaign, IL, 2024.
- [JN14a] Anatoli Juditski and Yuri Nesterov. Primal-dual subgradient methods for minimizing uniformly convex functions. *preprint arXiv:1401.1792*, 2014.

- [JN14b] Anatoli Juditsky and Yuri Nesterov. Deterministic and stochastic primal-dual subgradient algorithms for uniformly convex minimization. *Stochastic Systems*, 4(1):44–80, 2014.
- [JZ13] Rie Johnson and Tong Zhang. Accelerating stochastic gradient descent using predictive variance reduction. *Advances in neural information processing systems*, 26, 2013.
- [KF16] Donghwan Kim and Jeffrey A Fessler. Optimized first-order methods for smooth convex minimization. *Mathematical programming*, 159:81–107, 2016.
- [KF21] Donghwan Kim and Jeffrey A 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.
- [KGHN25] Yassine Kamri, François Glineur, Julien M Hendrickx, and Ion Necoara. On the worst-case analysis of cyclic block coordinate descent type algorithms. *arXiv preprint arXiv:2507.16675*, 2025.
- [KHG23] Yassine Kamri, Julien M Hendrickx, and François Glineur. On the worst-case analysis of cyclic coordinate-wise algorithms on smooth convex functions. In *2023 European Control Conference (ECC)*, pages 1–8. IEEE, IEEE, 2023.
- [Kim21] Donghwan Kim. Accelerated proximal point method for maximally monotone operators. *Mathematical Programming*, 190(1):57–87, 2021.
- [Kir34] Mojzesz Kirszbraun. Über die zusammenziehende und lipschitzsche transformationen. *Fundamenta Mathematicae*, 22(1):77–108, 1934.
- [KJH24] Qingyuan Kong, Rujun Jiang, and Yihan He. Linear convergence of the proximal gradient method for composite optimization under the polyak-łojasiewicz inequality and its variant. *arXiv preprint arXiv:2411.11628*, 2024.
- [KNS16] Hamed Karimi, Julie Nutini, and Mark Schmidt. Linear convergence of gradient and proximal-gradient methods under

- the Polyak-Lojasiewicz condition. In *Machine Learning and Knowledge Discovery in Databases: European Conference, Proceedings, Part I 16*, pages 795–811. Springer, 2016.
- [Kor76] Galina M Korpelevich. The extragradient method for finding saddle points and other problems. *Matecon*, 12:747–756, 1976.
- [KR20] Ahmed Khaled and Peter Richtárik. Better theory for sgd in the nonconvex world. *arXiv preprint arXiv:2002.03329*, 2020.
- [Las09] Jean Bernard Lasserre. *Moments, positive polynomials and their applications*, volume 1. World Scientific, 2009.
- [Lev44] Kenneth Levenberg. A method for the solution of certain non-linear problems in least squares. *Quarterly of applied mathematics*, 2(2):164–168, 1944.
- [LG24] Alan Luner and Benjamin Grimmer. Performance estimation for smooth and strongly convex sets. *preprint arXiv:2410.14811*, 2024.
- [Lie21] Felix Lieder. On the convergence rate of the halpern-iteration. *Optimization letters*, 15(2):405–418, 2021.
- [LMH15] Hongzhou Lin, Julien Mairal, and Zaid Harchaoui. A universal catalyst for first-order optimization. *Advances in neural information processing systems*, 28, 2015.
- [Loj63] Stanislaw Lojasiewicz. Une propriété topologique des sous-ensembles analytiques réels. *Les équations aux dérivées partielles*, 117:87–89, 1963.
- [LQT⁺24] Haochuan Li, Jian Qian, Yi Tian, Alexander Rakhlin, and Ali Jadbabaie. Convex and non-convex optimization under generalized smoothness. *Advances in Neural Information Processing Systems*, 36, 2024.
- [LRP16] 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.

- [LS67] Jacques-Louis Lions and Guido Stampacchia. Variational inequalities. *Communications on pure and applied mathematics*, 20(3):493–519, 1967.
- [LYR25] Jongmin Lee, Soheun Yi, and Ernest K Ryu. Convergence analyses of davis–yin splitting via scaled relative graphs. *SIAM Journal on Optimization*, 35(1):270–301, 2025.
- [Mar63] Donald W Marquardt. An algorithm for least-squares estimation of nonlinear parameters. *Journal of the society for Industrial and Applied Mathematics*, 11(2):431–441, 1963.
- [Mar08] Murray Marshall. *Positive Polynomials and Sums of Squares*, volume 146 of *Mathematical Surveys and Monographs*. American Mathematical Society, 2008.
- [MB11] Eric Moulines and Francis Bach. Non-asymptotic analysis of stochastic approximation algorithms for machine learning. *Advances in neural information processing systems*, 24, 2011.
- [MBB18] Siyuan Ma, Raef Bassily, and Mikhail Belkin. The power of interpolation: Understanding the effectiveness of sgd in modern over-parametrized learning. In *International Conference on Machine Learning*, pages 3325–3334. PMLR, 2018.
- [McS34] EJ McShane. Extension of range of functions. *Bulletin of the American Mathematical Society*, 40(12):837–842, 1934.
- [Mis23] Konstantin Mishchenko. Regularized Newton method with global $\mathcal{O}(1/k^2)$ convergence. *SIAM Journal on Optimization*, 33(3):1440–1462, 2023.
- [MJ20] Vien Mai and Mikael Johansson. Convergence of a stochastic gradient method with momentum for non-smooth non-convex optimization. In *International conference on machine learning*, pages 6630–6639. PMLR, 2020.
- [Mor65] Jean-Jacques Moreau. Proximité et dualité dans un espace hilbertien. *Bulletin de la Société mathématique de France*, 93:273–299, 1965.

- [MTB23] Céline Mouter, Adrien Taylor, and Francis Bach. A systematic approach to Lyapunov analyses of continuous-time models in convex optimization. *SIAM Journal on Optimization*, 33(3):1558–1586, 2023.
- [Nem94] Arkadi Nemirovski. Efficient methods in convex programming. *Lecture notes*, 1994.
- [Nes83] Yurii Nesterov. A method for solving the convex programming problem with convergence rate $o(1/k^2)$. In *Dokl akad nauk Sssr*, volume 269, page 543, 1983.
- [Nes08] Yu Nesterov. Accelerating the cubic regularization of Newton’s method on convex problems. *Mathematical Programming*, 112(1):159–181, 2008.
- [Nes12] Yu Nesterov. Efficiency of coordinate descent methods on huge-scale optimization problems. *SIAM Journal on Optimization*, 22(2):341–362, 2012.
- [Nes15] Yurii Nesterov. Universal gradient methods for convex optimization problems. *Mathematical Programming*, 152(1):381–404, 2015.
- [Nes18] Yurii Nesterov. *Lectures on convex optimization*, volume 137. Springer, 2018.
- [Nes21] Yurii Nesterov. Implementable tensor methods in unconstrained convex optimization. *Mathematical Programming*, 186(1):157–183, 2021.
- [Net19] Praneeth Netrapalli. Stochastic gradient descent and its variants in machine learning. *Journal of the Indian Institute of Science*, 99(2):201–213, 2019.
- [NJLS09] Arkadi Nemirovski, Anatoli Juditsky, Guanghui Lan, and Alexander Shapiro. Robust stochastic approximation approach to stochastic programming. *SIAM Journal on Optimization*, 19(4):1574–1609, 2009.
- [NLST17] Lam M Nguyen, Jie Liu, Katya Scheinberg, and Martin Takáč. Sarah: A novel method for machine learning problems using stochastic recursive gradient. In *International conference on machine learning*, pages 2613–2621. PMLR, 2017.

- [NN94] Yurii Nesterov and Arkadii Nemirovskii. Interior-point polynomial algorithms in convex programming. *SIAM*, 1994.
- [NP06] Yurii Nesterov and Boris T Polyak. Cubic regularization of Newton method and its global performance. *Mathematical programming*, 108(1):177–205, 2006.
- [NSW14] Deanna Needell, Nathan Srebro, and Rachel Ward. Stochastic gradient descent, weighted sampling, and the randomized kaczmarz algorithm. *Advances in neural information processing systems*, 27, 2014.
- [Nur73] Evgeni Alekseevich Nurminskii. The quasigradient method for the solving of the nonlinear programming problems. *Cybernetics*, 9(1):145–150, 1973.
- [Nur74] EA Nurminskii. Minimization of nondifferentiable functions in the presence of noise. *Cybernetics*, 10(4):619–621, 1974.
- [NY83] Arkadi Nemirovsky and David B. Yudin. *Problem Complexity and Method Efficiency in Optimization*. Wiley-Interscience, New York, 1983.
- [PJ92] Boris T Polyak and Anatoli B Juditsky. Acceleration of stochastic approximation by averaging. *SIAM journal on control and optimization*, 30(4):838–855, 1992.
- [Pol63] Boris T Polyak. Gradient methods for the minimisation of functionals. *USSR Computational Mathematics and Mathematical Physics*, 3(4):864–878, 1963.
- [Pol07] Boris T Polyak. Newton’s method and its use in optimization. *European Journal of Operational Research*, 181(3):1086–1096, 2007.
- [Pol09] Roman A Polyak. Regularized Newton method for unconstrained convex optimization. *Mathematical programming*, 120:125–145, 2009.
- [Pop80] Leonid Denisovich Popov. A modification of the arrow-hurwitz method of search for saddle points. *Mat. Zametki*, 28(5):777–784, 1980.

- [RB16] Ernest K Ryu and Stephen Boyd. Primer on monotone operator methods. *Appl. comput. math*, 15(1):3–43, 2016.
- [RBCH23] Anne Rubbens, Nizar Bousselmi, Sébastien Colla, and Julien M Hendrickx. Interpolation constraints for computing worst-case bounds in performance estimation problems. In *2023 62nd IEEE Conference on Decision and Control (CDC)*, pages 3015–3022. IEEE, 2023.
- [RBHG25] Anne Rubbens, Nizar Bousselmi, Julien M Hendrickx, and François Glineur. Performance estimation of second-order optimization methods on classes of univariate functions. *arXiv preprint arXiv:2506.22764*, 2025.
- [RCH25] Anne Rubbens, Sébastien Colla, and Julien M Hendrickx. Computer-aided analyses of stochastic first-order methods, via interpolation conditions for stochastic optimization. *arXiv preprint arXiv:2507.05466*, 2025.
- [RGP22] Teodor Rotaru, François Glineur, and Patrinos Panagiotis. Tight convergence rates of the gradient method on hypoconvex functions. *arXiv preprint arXiv:2203.00775*, 2022.
- [RGP24] 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.
- [RH24] Anne Rubbens and Julien M Hendrickx. A constraint-based approach to function interpolation, with application to performance estimation for weakly convex optimisation. *arXiv preprint arXiv:2405.08405*, 2024.
- [RHT25] Anne Rubbens, Julien M Hendrickx, and Adrien Taylor. A constructive approach to strengthen algebraic descriptions of function and operator classes. *arXiv preprint arXiv:2504.14377*, 2025.
- [RM51] Herbert Robbins and Sutton Monro. A stochastic approximation method. *The annals of mathematical statistics*, pages 400–407, 1951.

- [Roc70] Ralph Tyrrell Rockafellar. *Convex Analysis*. Princeton University Press, 1970.
- [RSB12] Nicolas Roux, Mark Schmidt, and Francis Bach. A stochastic gradient method with an exponential convergence rate for finite training sets. *Advances in neural information processing systems*, 25, 2012.
- [RSS11] Alexander Rakhlin, Ohad Shamir, and Karthik Sridharan. Making gradient descent optimal for strongly convex stochastic optimization. *arXiv preprint arXiv:1109.5647*, 2011.
- [RTBG20] Ernest K Ryu, Adrien B 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.
- [Rud76] Walter Rudin. *Principles of Mathematical Analysis*. McGraw-Hill, New York, 3rd edition, 1976.
- [Rud87] Walter Rudin. *Real and complex analysis*. McGraw-Hill, Inc., 1987.
- [Rup88] David Ruppert. Efficient estimations from a slowly convergent robbins-monro process. Technical report, Cornell University Operations Research and Industrial Engineering, 1988.
- [RVV20] Lorenzo Rosasco, Silvia Villa, and Bằng Công Vũ. Convergence of stochastic proximal gradient algorithm. *Applied Mathematics & Optimization*, 82:891–917, 2020.
- [RWW09] R.T. Rockafellar, M. Wets, and R.J.B. Wets. *Variational Analysis*. Grundlehren der mathematischen Wissenschaften. Springer, Berlin Heidelberg, 2009.
- [SGD21] Othmane Sebbouh, Robert M Gower, and Aaron Defazio. Almost sure convergence rates for stochastic gradient descent and stochastic heavy ball. In *Conference on Learning Theory*, pages 3935–3971. PMLR, 2021.
- [SL17] Ziqiang Shi and Rujie Liu. Better worst-case complexity analysis of the block coordinate descent method for large scale

- machine learning. In *2017 16th IEEE International Conference on Machine Learning and Applications (ICMLA)*, pages 889–892. IEEE, IEEE, 2017.
- [SLRB17] Mark Schmidt, Nicolas Le Roux, and Francis Bach. Minimizing finite sums with the stochastic average gradient. *Mathematical Programming*, 162:83–112, 2017.
- [SR13] Mark Schmidt and Nicolas Le Roux. Fast convergence of stochastic gradient descent under a strong growth condition. *arXiv preprint arXiv:1308.6370*, 2013.
- [SSZ13] Shai Shalev-Shwartz and Tong Zhang. Stochastic dual coordinate ascent methods for regularized loss. *The Journal of Machine Learning Research*, 14(1):567–599, 2013.
- [ST04] Daniel A Spielman and Shang-Hua Teng. Smoothed analysis of algorithms: Why the simplex algorithm usually takes polynomial time. *Journal of the ACM (JACM)*, 51(3):385–463, 2004.
- [Sta64] Guido Stampacchia. Formes bilineaires coercitives sur les ensembles convexes. *Comptes Rendus Hebdomadaires Des Seances De L Academie Des Sciences*, 258(18):4413, 1964.
- [STD19] Tianxiao Sun and Quoc Tran-Dinh. Generalized self-concordant functions: a recipe for Newton-type methods. *Mathematical Programming*, 178(1):145–213, 2019.
- [Sti19] Sebastian U Stich. Unified optimal analysis of the (stochastic) gradient method. *arXiv preprint arXiv:1907.04232*, 2019.
- [Sze39] Gabor Szeg. *Orthogonal polynomials*, volume 23. American Mathematical Soc., 1939.
- [TASV24] Cheik Traoré, Vassilis Apidopoulos, Saverio Salzo, and Silvia Villa. Variance reduction techniques for stochastic proximal point algorithms. *Journal of Optimization Theory and Applications*, pages 1–30, 2024.
- [Tay17] Adrien B Taylor. *Convex interpolation and performance estimation of first-order methods for convex optimization*. PhD thesis, Catholic University of Louvain, Louvain-la-Neuve, Belgium, 2017.

- [Tay24] Adrien Taylor. *Towards principled and systematic approaches to the analysis and design of optimization algorithms*. PSL Research University, 2024.
- [TB19] Adrien Taylor and Francis Bach. Stochastic first-order methods: non-asymptotic and computer-aided analyses via potential functions. In *Conference on Learning Theory*, pages 2934–2992. PMLR, 2019.
- [TD23] Adrien Taylor and Yoel Drori. An optimal gradient method for smooth strongly convex minimization. *Mathematical Programming*, 199(1):557–594, 2023.
- [THG17a] Adrien B. Taylor, Julien M. 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.
- [THG17b] Adrien B Taylor, Julien M 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. IEEE, 2017.
- [THG17c] Adrien B Taylor, Julien M Hendrickx, and François Glineur. Smooth strongly convex interpolation and exact worst-case performance of first-order methods. *Mathematical Programming*, 161:307–345, 2017.
- [THG18] Adrien B Taylor, Julien M Hendrickx, and François Glineur. Exact worst-case convergence rates of the proximal gradient method for composite convex minimization. *Journal of Optimization Theory and Applications*, 178:455–476, 2018.
- [Tie15] Heinrich Tietze. Über funktionen, die auf einer abgeschlossenen menge stetig sind. *Mathematische Annalen*, 77(1):1–22, 1915. doi:10.1007/BF01456923.
- [Toi24] Philippe L Toint. Examples of slow convergence for adaptive regularization optimization methods are not isolated. *arXiv preprint arXiv:2409.16047*, 2024.

- [TTD25] Daniel Berg Thomsen, Adrien Taylor, and Aymeric Dieuleveut. Tight analyses of first-order methods with error feedback, 2025. URL: <https://arxiv.org/abs/2506.05271>, arXiv:2506.05271.
- [TVSL18] Adrien Taylor, Bryan Van Scoy, and Laurent Lessard. Lyapunov functions for first-order methods: Tight automated convergence guarantees. In *International Conference on Machine Learning*, pages 4897–4906. PMLR, 2018.
- [ULG25] Manu Upadhyaya, Puya Latafat, and Pontus Giselsson. A lyapunov analysis of korpelevich’s extragradient method with fast and flexible extensions. *arXiv preprint arXiv:2502.00119*, 2025.
- [Val43] Frederick A Valentine. On the extension of a vector function so as to preserve a Lipschitz condition. *Bulletin of the American Mathematical Society*, 49:100–108, 1943.
- [Val45] Frederick Albert Valentine. A lipschitz condition preserving extension for a vector function. *American Journal of Mathematics*, 67(1):83–93, 1945.
- [VBS19] Sharan Vaswani, Francis Bach, and Mark Schmidt. Fast and faster convergence of sgd for over-parameterized models and an accelerated perceptron. In *The 22nd international conference on artificial intelligence and statistics*, pages 1195–1204. PMLR, 2019.
- [VG24] Pierre Vernimmen and François Glineur. Convergence analysis of an inexact gradient method on smooth convex functions. In *European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning*. ESANN, ESANN, 2024.
- [VG25a] Pierre Vernimmen and François Glineur. Empirical and computer-aided robustness analysis of long-step and accelerated methods in smooth convex optimization. *arXiv preprint arXiv:2506.09730*, 2025.
- [VG25b] Pierre Vernimmen and François Glineur. Worst-case convergence analysis of relatively inexact gradient descent on

- smooth convex functions. *arXiv preprint arXiv:2506.17145*, 2025.
- [VRN⁺24] Daniil Vankov, Anton Rodomanov, Angelia Nedich, Lalitha Sankar, and Sebastian U Stich. Optimizing (l_0, l_1) -smooth functions by gradient methods. *arXiv preprint arXiv:2410.10800*, 2024.
- [VSFL17] Bryan Van Scoy, Randy A Freeman, and Kevin M Lynch. The fastest known globally convergent first-order method for minimizing strongly convex functions. *IEEE Control Systems Letters*, 2(1):49–54, 2017.
- [VSL21] Bryan Van Scoy and Laurent Lessard. The speed-robustness trade-off for first-order methods with additive gradient noise. *arXiv preprint arXiv:2109.05059*, 2021.
- [Wel73] John C Wells. Differentiable functions on banach spaces with lipschitz derivatives. *Journal of Differential Geometry*, 8(1):135–152, 1973.
- [Whi34] Hassler Whitney. Differentiable functions defined in closed sets. I. *Transactions of the American Mathematical Society*, 36(2):369–387, 1934.
- [Wri15] Stephen J Wright. Coordinate descent algorithms. *Mathematical programming*, 151(1):3–34, 2015.
- [Xia09] Lin Xiao. Dual averaging method for regularized stochastic learning and online optimization. *Advances in Neural Information Processing Systems*, 22, 2009.
- [XZ14] Lin Xiao and Tong Zhang. A proximal stochastic gradient method with progressive variance reduction. *SIAM Journal on Optimization*, 24(4):2057–2075, 2014.
- [Yan14] Min Yan. Extension of convex function. *Journal of Convex Analysis*, 21(4):965–987, 2014.
- [ZAdK24] 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):243–276, 2024.

- [ZDS⁺19] Kaiwen Zhou, Qinghua Ding, Fanhua Shang, James Cheng, Danli Li, and Zhi-Quan Luo. Direct acceleration of saga using sampled negative momentum. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 1602–1610. PMLR, 2019.
- [ZG23] Moslem Zamani and François Glineur. Exact convergence rate of the last iterate in subgradient methods. *arXiv preprint arXiv:2307.11134*, 2023.
- [ZHSJ19] Jingzhao Zhang, Tianxing He, Suvrit Sra, and Ali Jadbabaie. Why gradient clipping accelerates training: A theoretical justification for adaptivity. In *International Conference on Learning Representations*, 2019.
- [ZZ15] Peilin Zhao and Tong Zhang. Stochastic optimization with importance sampling for regularized loss minimization. In *Proceedings of Machine Learning Research*, volume 37, pages 1–9. PMLR, 2015.