
Linear Lower Bounds and Conditioning of Differentiable Games

Adam Ibrahim¹ Waïss Azizian² Gauthier Gidel¹ Ioannis Mitliagkas¹

Abstract

Recent successes of game-theoretic formulations in ML have caused a resurgence of research interest in differentiable games. Overwhelmingly, that research focuses on methods and upper bounds on their speed of convergence. In this work, we approach the question of fundamental iteration complexity by providing lower bounds to complement the linear (i.e. geometric) upper bounds observed in the literature on a wide class of problems. We cast saddle-point and min-max problems as 2-player games. We leverage tools from single-objective convex optimisation to propose new linear lower bounds for convex-concave games. Notably, we give a linear lower bound for n -player differentiable games, by using the spectral properties of the update operator. We then propose a new definition of the condition number arising from our lower bound analysis. Unlike past definitions, our condition number captures the fact that linear rates are possible in games, even in the absence of strong convexity or strong concavity in the variables.

1. Introduction

Game formulations arise commonly in many fields, such as game theory (Harker and Pang, 1990), machine learning (Kim and Boyd, 2008; Goodfellow et al., 2014), and computer vision (Chambolle and Pock, 2011; Wang et al., 2014) among others, and encompass saddle-point problems (Palaaniappan and Bach, 2016; Chambolle and Pock, 2011; Chen et al., 2017).

The machine learning community has been overwhelmingly using gradient-based methods to train differentiable games (Goodfellow et al., 2014; Salimans et al., 2016). These methods are not designed with game dynamics in

mind (Mescheder et al., 2017), and to make matters worse, have been often tuned suboptimally (Gidel et al., 2019b). A recent series of publications in machine learning brings in tools from the minimax and game theory literature to offer better, faster alternatives (Daskalakis et al., 2018; Gidel et al., 2019a;b). This exciting trend begs the question: how fast can we go? Knowing the fundamental limits of this class of problems is critical in steering future algorithmic research.

In order to answer this question, the optimisation literature contains a few different approaches based on the distance between the iterates at a step t and the optimal choice of parameters. Given an optimisation algorithm, it is possible to show under certain assumptions on the objectives that this error is in $\mathcal{O}(\rho^t)$, where the rate of convergence ρ depends on the algorithm (Nesterov, 2004). If $\rho \in (0, 1)$, we say that the rate of convergence is linear, which corresponds to the error decaying exponentially fast. Hence, a lower bound on the rate of convergence limits how fast an algorithm may converge. This is important as it helps establish the tightness of upper bounds, which happens when they are matched by the lower bounds, and may otherwise indicate possible acceleration of the method considered. For example, in single-objective optimisation, the lower bound on the rate of convergence of first-order black box algorithms is known to be linear for smooth, strongly convex objectives, and can be derived via a domino-like coverage argument by Nesterov (2004). Another recent, spectral approach by Arjevani et al. (2016) complement these results by proposing linear lower bounds for a large class of optimisers in finite-dimensional settings. As the lower bounds for Nesterov’s accelerated gradient obtained by those techniques match the upper bound, we know that Nesterov’s accelerated gradient is optimal within a large class of methods for smooth, strongly convex objectives. Additionally, in optimisation, a natural concept of condition number arises to describe the difficulty of μ -strongly convex, L -smooth objectives. This condition number is the only problem-dependent quantity that appears in both the upper and lower bounds and is given by $\kappa = L/\mu$ (Nesterov, 2004). In optimisation, there is a clear distinction between strongly convex objectives, where the condition number is finite and linear rates are achievable, and general convex objectives where the condition number can be undefined and only sublinear rates are possible in

*Equal contribution ¹Mila, University of Montreal ²Ecole Normale Supérieure, Paris. Correspondence to: Adam Ibrahim <<first>.<last>@umontreal.ca>.

general.

When studying lower bounds for convex-concave min-max, one is faced with a number of distinct challenges compared to the optimisation setting. In particular, there is no universally accepted definition of a condition number. Some commonly used definitions, like the one used in [Chambolle and Pock \(2011\)](#); [Palaniappan and Bach \(2016\)](#), are undefined for bilinear problems, which lack strong convexity and strong concavity in the variables. This is problematic because we know that both extragradient and gradient methods with negative momentum achieve linear convergence in bilinear games ([Korpelevich, 1976](#); [Gidel et al., 2019b](#)). *Can we get a condition number that captures the fact that linear rates are possible even in the absence of strong convexity and strong concavity?*

We show that it is possible by providing new lower bounds, obtained by casting saddle-point and min-max problems as games and leveraging existing proof techniques originally designed for smooth strongly convex, single-objective optimisation. These bounds also yield a meaningful condition number for the bilinear case, in the absence of strong convexity and strong concavity in the variables. Our contributions are summarised as follows:

1. We generalise *Nesterov’s domino argument* and design a difficult min-max problem to derive a linear lower bound on the rate of convergence of several first-order black box optimisation algorithms for 2-player games and min-max problems. In order to get an asymptotic rate using the domino bound, one needs to resort to the analysis of infinite-dimensional problems.
2. We propose a linear lower bound for finite-dimensional problems by generalising the p -SCLI framework proposed by [Arjevani et al. \(2016\)](#) to n -objective optimisation algorithms. This lower bound stems from the spectral properties of the algorithms on quadratics, and is valid for any number of players, and in particular 2-player games and min-max problems. This bound is tight for $n = 1$ since it reduces to the one presented by [Arjevani et al. \(2016\)](#) for strongly convex, smooth single-player optimisation.
3. We provide a formulation of the condition number of 2-player games consistent with the existing literature on upper bounds for games and min-max problems. In particular, this condition number is finite for bilinear games.

After the results of this work were made available online, several researchers have proposed methods to match some of our bounds in the smooth strongly-convex-strongly-concave setting ([Fallah et al., 2020](#); [Lin et al., 2020](#)) and the bilinear setting ([Azizian et al., 2020](#)), which is merely

convex-concave, thereby establishing the tightness of some of the bounds and the optimality of those methods in those regimes.

The rest of the paper is organised as follows. We purposely discuss preliminaries first in Section 2 to introduce the general framework used to present in Section 3 the relevant literature in the context of our results. In Section 4, we provide lower bounds using Nesterov’s domino argument, and in Section 5 we improve on those bounds using the spectral technique. We conclude with some discussion.

2. Preliminaries

2.1. Differentiable games

Following the definition of [Balduzzi et al. \(2018\)](#), a *differentiable game* is characterised by n players, each associated with a set of parameters $\mathbf{w}_i \in \mathbb{R}^{d_i}$ and a twice continuously differentiable objective function $l_i : \mathbb{R}^d \rightarrow \mathbb{R}$ of all the parameters $\mathbf{w} = (\mathbf{w}_1, \dots, \mathbf{w}_n) \in \mathbb{R}^d$, where $d = \sum_{i=1}^n d_i$. In particular, if $\sum_{i=1}^n l_i(\mathbf{w}) = 0$, we say that the game is *zero-sum*.

Often, we seek to minimise the objectives l_i , and look for *Nash equilibria* $\mathbf{w}^* = (\mathbf{w}_1^*, \dots, \mathbf{w}_n^*)$, which satisfy¹ for all i

$$\mathbf{w}_i^* \in \arg \min_{\mathbf{w}_i} l_i(\mathbf{w}_1^*, \dots, \mathbf{w}_{i-1}^*, \mathbf{w}_i, \mathbf{w}_{i+1}^*, \dots, \mathbf{w}_n^*). \quad (1)$$

In order to find the Nash equilibria, we may look for stationary points, corresponding to the zeros of the vector field $\mathbf{v}(\mathbf{w}) = (\nabla_{\mathbf{w}_1} l_1(\mathbf{w}) \dots \nabla_{\mathbf{w}_n} l_n(\mathbf{w}))^\top$. In single-objective optimisation, which corresponds to a 1-player game, we know that stationary points of \mathbf{v} do not necessarily represent minima of the objective function, and higher order information, such as the Hessian, is necessary to determine whether a stationary point is a minimum. The same is true for a game with several players ([Balduzzi et al., 2018](#)), where the *Jacobian* of \mathbf{v} , given by

$$\nabla \mathbf{v}(\mathbf{w}) = \begin{pmatrix} \nabla_{\mathbf{w}_1}^2 l_1(\mathbf{w}) & \dots & \nabla_{\mathbf{w}_n} \nabla_{\mathbf{w}_1} l_1(\mathbf{w}) \\ \vdots & & \vdots \\ \nabla_{\mathbf{w}_1} \nabla_{\mathbf{w}_n} l_n(\mathbf{w}) & \dots & \nabla_{\mathbf{w}_n}^2 l_n(\mathbf{w}) \end{pmatrix} \quad (2)$$

gives sufficient conditions to determine whether a stationary point is a Nash equilibrium. Note that our lower bound analysis encompasses games with stable stationary points that are not Nash equilibria.

¹Of course, we could be trying to maximise some players’ objectives, but we can without loss of generality work with minima since $\arg \max f = \arg \min(-f)$

2.2. Quadratic games

In order to gain insight on general games, we focus on quadratic games², corresponding to games with quadratic objectives l_i . In our analysis, we will mostly discuss two-player games, where the players respectively control the parameters $\mathbf{x} \in \mathbb{R}^{d_1}$ and $\mathbf{y} \in \mathbb{R}^{d_2}$. The quadratic objectives take the form

$$\begin{aligned} l_1(\mathbf{x}, \mathbf{y}) &= \frac{1}{2} \mathbf{x}^\top \mathbf{S}_1 \mathbf{x} + \mathbf{x}^\top \mathbf{M}_{12} \mathbf{y} + \mathbf{x}^\top \mathbf{b}_1 \\ l_2(\mathbf{x}, \mathbf{y}) &= \frac{1}{2} \mathbf{y}^\top \mathbf{S}_2 \mathbf{y} + \mathbf{y}^\top \mathbf{M}_{21} \mathbf{x} + \mathbf{y}^\top \mathbf{b}_2 \end{aligned} \quad (3)$$

with \mathbf{S}_1 and \mathbf{S}_2 symmetric. In that case the vector field is given by

$$\begin{aligned} \mathbf{v}(\mathbf{x}, \mathbf{y}) &= \begin{pmatrix} \mathbf{S}_1 \mathbf{x} + \mathbf{M}_{12} \mathbf{y} + \mathbf{b}_1 \\ \mathbf{M}_{21} \mathbf{x} + \mathbf{S}_2 \mathbf{y} + \mathbf{b}_2 \end{pmatrix} \\ &= \mathbf{A} \begin{pmatrix} \mathbf{x} \\ \mathbf{y} \end{pmatrix} + \mathbf{b} \end{aligned} \quad (4)$$

$$\text{with } \mathbf{A} \triangleq \begin{pmatrix} \mathbf{S}_1 & \mathbf{M}_{12} \\ \mathbf{M}_{21} & \mathbf{S}_2 \end{pmatrix}, \quad \mathbf{b} \triangleq \begin{pmatrix} \mathbf{b}_1 \\ \mathbf{b}_2 \end{pmatrix}$$

where \mathbf{A} is the Jacobian of \mathbf{v} . For n -player quadratic games, the vector field and Jacobian \mathbf{A} take the form

$$\begin{aligned} \mathbf{v}(\mathbf{w}) &= \mathbf{A} \mathbf{w} + \mathbf{b} \\ \text{with } \mathbf{A} &\triangleq \begin{pmatrix} \mathbf{S}_1 & \dots & \mathbf{M}_{1n} \\ \vdots & & \vdots \\ \mathbf{M}_{n1} & \dots & \mathbf{S}_n \end{pmatrix}, \quad \mathbf{b} \triangleq \begin{pmatrix} \mathbf{b}_1 \\ \vdots \\ \mathbf{b}_n \end{pmatrix} \end{aligned} \quad (5)$$

where the \mathbf{S}_i are symmetric. For more details on quadratic games, see Appendix A. Further assumptions on the dimensionality or properties of \mathbf{S}_i will be introduced in the rest of the paper as they become relevant (e.g. positive semi-definiteness in the next subsection).

We shall henceforth refer to the Jacobian of the vector field of quadratic n -player games simply as the Jacobian, since our analysis will be solely based on quadratic objectives. Interestingly, the problem of finding \mathbf{w} such that $\mathbf{v}(\mathbf{w}) = 0$ consists in solving a system of linear equations (SLE) (Richardson, 1911). In fact, several techniques to precondition systems of linear equation make use of casting the SLE as a game and optimising it with proximal methods, such as Benzi and Golub (2004).

In this paper, we will denote the spectrum of a matrix \mathbf{M} by $\sigma(\mathbf{M})$, and define the *block spectral bounds* $\mu_1, \mu_2, \mu_{12}, L_1, L_2, L_{12}$ as constants bounding the spectra of the blocks in the Jacobian of eq. 4:

$$\begin{aligned} \mu_1 &\leq |\sigma(\mathbf{S}_1)| \leq L_1 & \mu_2 &\leq |\sigma(\mathbf{S}_2)| \leq L_2 \\ \mu_{12}^2 &\leq |\sigma(\mathbf{M}_{12} \mathbf{M}_{12}^\top)| \leq L_{12}^2 \end{aligned} \quad (6)$$

²Note that quadratic games are inherently relevant; e.g. in reinforcement learning to learn a linear value function from the mean squared projected Bellman error (Du et al., 2017)

where we assume that \mathbf{M}_{12} is a wide or square matrix (if it is a tall matrix, we use $\mu_{12}^2 \leq |\sigma(\mathbf{M}_{12}^\top \mathbf{M}_{12})| \leq L_{12}^2$ to define μ_{12} and L_{12} instead of the last inequality of eq. 6).

2.3. Min-max of quadratics as 2-player quadratic games

Consider the family \mathcal{P} of min-max problems of the form

$$\begin{aligned} \min_{\mathbf{x} \in \mathbb{R}^{d_1}} \max_{\mathbf{y} \in \mathbb{R}^{d_2}} f(\mathbf{x}, \mathbf{y}) &= \mathbf{x}^\top \mathbf{M} \mathbf{y} + \frac{1}{2} \mathbf{x}^\top \mathbf{S}_1 \mathbf{x} - \frac{1}{2} \mathbf{y}^\top \mathbf{S}_2 \mathbf{y} \\ &\quad + \mathbf{x}^\top \mathbf{b}_1 - \mathbf{y}^\top \mathbf{b}_2 + c \end{aligned} \quad (\mathcal{P})$$

where $\sigma(\mathbf{M} \mathbf{M}^\top), \sigma(\mathbf{S}_1), \sigma(\mathbf{S}_2) \subseteq [0, +\infty)$

with possible constraints and where the dimension need not be finite, e.g. $\mathbf{x}, \mathbf{y} \in \ell_2 \triangleq \{\mathbf{u} \in \mathbb{R}^N \mid \sum_i^\infty \mathbf{u}_i^2 < \infty\}$. The optimisation of such a problem is equivalent to finding a pair $(\mathbf{x}^*, \mathbf{y}^*)$ such that,

$$\mathbf{x}^* \in \arg \min f(\mathbf{x}, \mathbf{y}^*) \quad \text{and} \quad \mathbf{y}^* \in \arg \max f(\mathbf{x}^*, \mathbf{y}) \quad (7)$$

Noting that $\arg \max f = \arg \min(-f)$, we get that this optimisation problem is equivalent to a zero-sum 2-player game with objectives $l_1 = -l_2 = f$ (see eq. 1). This problem can be reduced to searching for the Nash equilibria of the 2-player quadratic game with simplified objectives $l_x(\mathbf{x}, \mathbf{y}) = \frac{1}{2} \mathbf{x}^\top \mathbf{S}_1 \mathbf{x} + \mathbf{x}^\top \mathbf{M} \mathbf{y} + \mathbf{x}^\top \mathbf{b}_1$ and $l_y(\mathbf{x}, \mathbf{y}) = \frac{1}{2} \mathbf{y}^\top \mathbf{S}_2 \mathbf{y} - \mathbf{x}^\top \mathbf{M} \mathbf{y} + \mathbf{y}^\top \mathbf{b}_2$, where the \mathbf{S}_i have been symmetrised (see Appendix A for an explanation of the symmetrisation of \mathbf{S}_i and why the objectives can be simplified). Eq. 4 yields the vector field

$$\mathbf{v}(\mathbf{x}, \mathbf{y}) = \begin{pmatrix} \mathbf{S}_1 & \mathbf{M} \\ -\mathbf{M}^\top & \mathbf{S}_2 \end{pmatrix} \begin{pmatrix} \mathbf{x} \\ \mathbf{y} \end{pmatrix} + \begin{pmatrix} \mathbf{b}_1 \\ \mathbf{b}_2 \end{pmatrix} \quad (8)$$

Therefore, the pair $(\mathbf{x}^*, \mathbf{y}^*)$ from eq. 7 exists if and only if the corresponding games with vector field given in eq. 8 admit a Nash equilibrium since $\mathbf{S}_1, \mathbf{S}_2 \succeq 0$. Note that we could also go from a quadratic game satisfying the above to a min-max formulation. Hence, any lower bound on quadratic games of the form of eq. 8 is a lower bound on min-max problems (in \mathcal{P}), and vice versa.

3. Background

3.1. Existing bounds for 2-player quadratic min-max problems

Some upper bounds on the rate of convergence of certain optimisation algorithms exist for unconstrained problems in \mathcal{P} . These upper bounds on the rate of convergence ρ imply that for any problem in \mathcal{P} , the iterates will converge to a solution in $\mathcal{O}(\rho^t)$. For clarity's sake, we reformulate these upper bounds to be consistent with the notation of \mathcal{P} and eq. 6. Letting $\kappa = \frac{L_{12}}{\sqrt{\mu_1 \mu_2}}$, Chen and Rockafellar (1997) analyse the forward-backward algorithm, and

find a convergence in $\mathcal{O}\left(\left(\sqrt{1 - \left(\frac{\min(\mu_1, \mu_2, \mu_{12})}{\max(L_1, L_2, L_{12})}\right)^2}\right)^t\right)$.

Chambolle and Pock (2011) give an algorithm for which the convergence is in $\mathcal{O}\left(\left(\sqrt{1 - \frac{2}{\kappa+2}}\right)^t\right)$. Palaniappan and Bach (2016) present an accelerated version of the forward-backward algorithm with variance reduction with convergence in $\mathcal{O}\left(\left(1 - \frac{1}{1+2\kappa}\right)^t\right)$. Note that asymptotically, the Chambolle-Pock and accelerated Forward-Backward rates match up to a factor of 2 on κ . Finally, Gidel et al. (2019b) give an upper bound in $\mathcal{O}\left(\left(1 - \frac{1}{4L_{12}^2/\mu_{12}^2}\right)^t\right)$ on the convergence of alternating gradient descent with negative momentum for non-singular *bilinear games*, i.e. quadratic games satisfying eq. 8 with $\mathbf{S}_1 = \mathbf{S}_2 = 0$ and non-singular Jacobian.

A key problem with those rates of convergence is that the tightness of the upper bound is not established. Such information is important since it may indicate that the algorithm can be accelerated. Ideally, one would use the rate of convergence of the hardest problem (i.e. slowest convergence) in the class of problems, which would be a tight upper bound. If one can only find a lower bound on the rate of convergence of the hardest problem, then any upper bound on the entire problem class must be greater than that lower bound to avoid a contradiction. This is because the (upper bound on the) rate of convergence for a class of problems must apply to *any* problem in the class, and hence be greater than any lower bound derived on any particular problem within that class. Usually, it is not possible to find the problem with the slowest convergence, so one may have to guess a hard enough problem. If the lower bound on that problem matches the upper bound for the problem class, then we have established that not only this problem is one of the hardest problems in the class, but also that the upper bound is tight. Therefore, it is important to remember that the goal of lower bounds generally is not to apply to every problem within the class, unlike upper bounds, but to help estimate how much the upper bounds on the whole class can be improved given the presence of hard problems in the class.

Unlike upper bounds, relevant lower bounds for first-order methods on saddle-point problems are scarcer in the literature. Nemirovsky (1992) gives a lower bound in $\mathcal{O}(1/t)$ for a limited number of steps. Ouyang and Xu (2018) also leverage Krylov subspace techniques, and show lower bounds in $\mathcal{O}(1/t)$ in the monotone case and $\mathcal{O}(1/t^2)$ in the strongly monotone case, assuming the number of iterations is less than half the dimension of the parameters. Note that Ouyang and Xu (2018) do not assume smoothness of the objective in \mathbf{y} . A key issue is that since these bounds are only valid for

a limited number of steps, they do not yield bounds that can be compared with the upper bounds previously mentioned. In contrast, the lower bounds presented in this work are valid for any number of steps and are linear, and therefore provide a direct limit to the acceleration of methods achieving linear convergence on two-player games. Additionally, our lower bounds also yield condition numbers that give intuition about the difficulty inherent to a problem, and can be computed in a plug-and-play fashion using either bounds on the spectrum of the full Jacobian, or on the spectra of its blocks.

3.2. Lower bound techniques for convex optimisation with bounded spectrum

In single-objective optimisation, i.e. a 1-player game, the Jacobian in eq. 2 reduces to the Hessian of the objective, denoted $\mathbf{H}(\mathbf{x})$. In that case, if there exists $\mu, L \in \mathbb{R}^{++}$ such that for all \mathbf{x} in the domain considered

$$\mu \preceq \mathbf{H}(\mathbf{x}) \preceq L$$

the objective is μ -strongly convex and has L -Lipschitz gradients, and the convergence rates (i.e. upper bounds) are known to be linear in the number of iterations for several classes of algorithms (Nemirovsky and Yudin, 1983; Nesterov, 2004). In the context of convex minimisation, various lower bounds have been derived depending on whether the objective is strongly convex and/or has Lipschitz gradients (see (Bubeck et al., 2015) for an overview).

Nesterov’s lower bound In particular, Nesterov (2004) gives an information-based complexity bound for μ -strongly convex objectives with L -Lipschitz gradients, by showing that there is a μ -strongly convex example in $\ell_2 \rightarrow \mathbb{R}$ with L -Lipschitz gradients for which first-order black box methods, i.e. methods using only past iterates and gradients of past iterates at every update, converge linearly at a rate at least $\rho = 1 - \frac{2}{\sqrt{\kappa+1}}$, where the condition number is given by $\kappa = L/\mu$. The proof relies on the fact that at iteration t , only the t first components of the estimates \mathbf{x}_t have been updated from their initial values, where $\mathbf{x}_0 = 0$. This is then used to lower bound the distance to the optimum. Since an infinite number of iterations is required to converge in ℓ_2 if the solution \mathbf{x}^* has an infinite number of nonzero components, we obtain asymptotic rates. An important caveat is that an infinite-dimensional example does not directly yield a lower bound for finite-dimensional problems.

p -SCLI Arjevani et al. (2016) introduce the p -SCLI framework to provide bounds for a large class of methods used for optimising μ -strongly convex objectives with L -Lipschitz gradients. Roughly speaking, an algorithm is p -SCLI if its update rule on quadratics $f(\mathbf{x}) = \mathbf{x}^\top \mathbf{A} \mathbf{x} + \mathbf{x}^\top \mathbf{b}$ with \mathbf{A} symmetric is a linear combination of the p previous iterates and \mathbf{b} , where the coefficients are matrices that depend on

\mathbf{A} and are assumed to be simultaneously triangularisable. The spectral properties of the update rule are used to derive lower bounds on the rate of convergence of p -SCLI algorithms. The lower bound on the rate of convergence of p -SCLI methods is given by $\rho = 1 - \frac{2}{\sqrt[\kappa]{\kappa+1}}$ for $\kappa = L/\mu$. This allows us to recover lower bounds for gradient descent ($p = 1$) or Nesterov (1983)'s accelerated gradient descent ($p = 2$) that match the upper bounds. A key advantage is that these bounds are more refined than Nesterov's by introducing the dependence on p (for example, both GD and Nesterov's accelerated gradient descent are black box first-order methods, but this bound is tighter for $p = 1$ methods such as GD), and do not rely on an infinite-dimensional example, but rather on the spectral properties of the methods. Such bounds highlight how lower bounds can suggest potential acceleration: in the example previously given, the addition of momentum to gradient descent turns the method from $p = 1$ to $p = 2$, thereby decreasing the lower bound and allowing for potentially faster convergence. Finally, the p -SCLI framework also yields upper bounds, and the authors also show a general method to accelerate algorithms on quadratics with the hope that the acceleration is relevant to more general classes of objectives, albeit at a cost too prohibitive to be practical.

Interestingly, both techniques produce a tight lower bound from a quadratic objective, indicating that quadratics are asymptotically as hard as any other strongly convex, smooth problem in single-objective optimisation. This motivates the use of quadratic games to derive the lower bounds presented in this paper as we generalise those methods to the multi-objective setting. When there are several players, however, the Jacobian is no longer symmetric, and its spectrum will generally be complex, and hence several of the arguments used in single-objective optimisation fail to apply directly.

4. Parametric Lower Bounds from Nesterov's Domino Argument

In this section, we will only discuss min-max problems. The class \mathcal{F} of counterexamples considered is

$$\min_{\mathbf{x} \in \ell_2} \max_{\mathbf{y} \in \ell_2} f(\mathbf{x}, \mathbf{y}) = c\mathbf{x}^\top \mathbf{M}\mathbf{y} - d_1\mathbf{x}^\top \mathbf{e}_1 + d_2\mathbf{y}^\top \mathbf{e}_1 + \frac{\mu_1}{2} \|\mathbf{x}\|^2 - \frac{\mu_2}{2} \|\mathbf{y}\|^2 \quad (\mathcal{F})$$

where \mathbf{M} is an infinite-dimensional bidiagonal matrix i.e. $\forall i, M_{ii} = a_0$ and $M_{i,i+1} = a_1$ with all other entries set to 0, such that $ca_0a_1 \neq 0$ and $\mu_1, \mu_2 \in \mathbb{R}^{++}$. Since these problems are in \mathcal{P} , the lower bounds of this section are in particular bounds on the optimisation of min-max problems.

Definition 1 (Two-step linear span assumption). *A first-order black box method for 2-player games satisfies the two-step linear span assumption on \mathcal{F} if for problems in \mathcal{F}*

with Jacobian \mathbf{A} (cf eq. 5):

$$\mathbf{w}_t \in \mathbf{w}_0 + \text{Span}(\mathbf{w}_0, \dots, \mathbf{w}_{t-1}, \mathbf{A}\mathbf{w}_0, \dots, \mathbf{A}\mathbf{w}_{t-1}, \mathbf{A}^2\mathbf{w}_0, \dots, \mathbf{A}^2\mathbf{w}_{t-1}, \mathbf{b}, \mathbf{A}\mathbf{b}) \quad (9)$$

Examples of such methods include simultaneous gradient descent, negative momentum and extragradient. One way to design challenging problems for these methods is to construct problems with a dense solution $(\mathbf{x}^*, \mathbf{y}^*)$ for which only one new component of the iterates may change from its initial value at every iteration (Nesterov, 2004), a phenomenon we will refer to as the *domino argument* (see Appendix B.1 for some intuition, where we show that the argument also applies to cases where diagonal matrices are used as coefficients in the span, and to alternating implementations of any algorithm satisfying the two-step linear span assumption on \mathcal{F} thanks to the properties of bidiagonal Toeplitz matrices).

4.1. A first lower bound for games with block spectral bounds μ_1, μ_2, L_{12}

Proposition 2 (Naive bound). *For any problem class containing quadratic games, there exists a function $f : \ell_2 \times \ell_2 \rightarrow \mathbb{R}$ corresponding to a problem in \mathcal{P} with block spectral bounds $\mu_1 = L_1, \mu_2 = L_2, L_{12} \in \mathbb{R}^{++}$ as defined in eq. 6, that has condition number $\kappa = \frac{L_{12}}{\sqrt{\mu_1\mu_2}}$ such that for any number of iterations $t \geq 1$ and any procedure satisfying the two-step linear span assumption (see def. 1), the following lower bound holds:*

$$\|(\mathbf{x}_t, \mathbf{y}_t) - (\mathbf{x}^*, \mathbf{y}^*)\| \geq \left(1 - \frac{2}{\sqrt{\kappa^2 + 1} + 1}\right)^{t+1} \cdot \|(\mathbf{x}_0, \mathbf{y}_0) - (\mathbf{x}^*, \mathbf{y}^*)\| \quad (10)$$

We invite the reader to consult Appendix B.2 for the proof. This lower bound on the distance to a solution also yields a lower bound on the minimum number of steps necessary for all subsequent iterates to be within a target distance — typically referred to as *iteration complexity*. We may interpret the proposition as being a bound for 2-player games with block spectral bounds μ_1, μ_2 , and L_{12} as the bound provided holds for a problem sharing the same block spectral bounds. Similarly, we also get a bound on problems with block spectral bounds L_1, L_2 and L_{12} by replacing μ_i by L_i appropriately in κ .

We appear to obtain the same condition number as in the upper bound literature. If we assume this bound and condition number to be representative of a finite-dimensional bound as was the case in convex optimisation, we easily see an apparent contradiction from the upper bound on the rate of convergence of alternating gradient descent with negative momentum for bilinear games given by (Gidel et al.,

2019b). Indeed, if we let $\mu_1, \mu_2 \rightarrow 0$, the rate of convergence in Prop. 2 goes to 1, whereas the upper bound of negative momentum is not affected and may indicate fast convergence. This illustrates how the condition number of the upper bounds is not general enough to be representative of inherent difficulty: it can be shown that for the problem used in the proof of the proposition, $\mu_{12} = 0$. As such, it is not surprising that the bound failed to hold against the upper bound of negative momentum on bilinear games; they were not comparable as by definition bilinear games have $\mu_{12} > 0$. This shows that μ_{12} encodes critical information that this condition number was not able to capture. Nevertheless, an important point is that the bound itself is correct and represents a problem with slow convergence; it just fails to yield a condition number that accurately captures difficulty as μ_{12} does not appear.

However, by refining our proof technique, we can derive a bound which avoids this issue, and yields tighter bounds for games for which we know $\mu_{12}, L_{12}, \mu_1, \mu_2$.

4.2. Improved lower bound for games with block spectral bounds $\mu_1, \mu_2, \mu_{12}, L_{12}$

Theorem 3. *For any problem class containing quadratic games, there exists a function $f : \ell_2 \times \ell_2 \rightarrow \mathbb{R}$ corresponding to a problem in \mathcal{P} with block spectral bounds $\mu_1 = L_1, \mu_2 = L_2, L_{12} \in \mathbb{R}^{++}, \mu_{12} \in \mathbb{R}^+$ as defined in eq. 6, that has condition number $\kappa = \sqrt{\frac{L_{12}^2 + \mu_1 \mu_2}{\mu_{12}^2 + \mu_1 \mu_2}}$, such that for any number of iterations $t \geq 1$ and any procedure satisfying the two-step linear span assumption, the following lower bound holds:*

$$\|(\mathbf{x}_t, \mathbf{y}_t) - (\mathbf{x}^*, \mathbf{y}^*)\| \geq \left(1 - \frac{2}{\kappa + 1}\right)^{t+1} \cdot \|(\mathbf{x}_0, \mathbf{y}_0) - (\mathbf{x}^*, \mathbf{y}^*)\| \quad (11)$$

The same result holds for any problem class containing bilinear games, if one sets $\mu_1 = L_1 = \mu_2 = L_2 = 0$.

Corollary 4 (Iteration complexity bound). *For the same problem classes and under the same assumptions as Theorem 3, the minimal number of steps t required to reach a target distance ϵ from the solution, that is $\|(\mathbf{x}_t, \mathbf{y}_t) - (\mathbf{x}^*, \mathbf{y}^*)\| < \epsilon$, is given by*

$$t \geq \frac{\kappa - 1}{2} \log \left(\frac{\|(\mathbf{x}_0, \mathbf{y}_0) - (\mathbf{x}^*, \mathbf{y}^*)\|}{\epsilon} \right) - 1 \quad (12)$$

This generalises Prop. 2. The proof can be found in Appendix B.3, where we also show how we used spectral properties of Toeplitz matrices in Banach algebras to create the hard problems yielding the bound. The proof of the iteration complexity bound can be found in Appendix D. As with the lower bounds, note that the iteration complexity bound does not necessarily hold for every problem in the class; the

idea is that to reach a target error ϵ we know that there is at least one problem that requires at least the number of steps indicated in the bound, and therefore that to optimise over a problem class that satisfies the assumptions, we will need in general at least the number of steps given in the bound, to account for the hard problems.

It is important to emphasize that as is the case with Nesterov’s argument for single-objective optimisation, this bound is still based on an infinite-dimensional problem, and that upper bounds generally are proven for finite-dimensional settings. We may hope that this bound also holds in finite dimension, since we are not aware of upper bounds contradicting it, and were not able to generate finite-dimensional 2-player games for which the bound did not hold empirically. The condition number appearing in Thm. 3 is more expressive than the one found in the upper bound literature $\kappa = L_{12}/\sqrt{\mu_1 \mu_2}$, and is lower bounded by 1, instead of 0. A limitation, however, is that this κ is not able to dissociate $L_1 \neq \mu_1, L_2 \neq \mu_2$, which is a problem in terms of expressivity of the condition number. It may also threaten the tightness of the lower bound since intuition from convex optimisation would suggest that objectives with matching lower and upper bounds on the spectra are easier to optimise. This stems from the fact that the closed form solution for problems in \mathcal{P} when \mathbf{S}_i is non-scalar, which we would need for all block spectral bounds to appear in κ , is complicated and the associated condition number is impractical. Therefore, we leave the matter of deriving a practical bound based on the domino argument involving all μ_i and L_i as future work.

Interestingly, the rate takes the same form as in strongly convex smooth optimisation, suggesting that for general n -player games, we may still get a lower bound of the form $\rho \geq 1 - \frac{2}{\kappa+1}$ for some generalised condition number κ . This intuition will be highlighted in the next section, by deriving lower bounds from the spectral properties of the update operators of a large class of optimisation methods for n -player games. The results we are about to introduce will also address the matter of $L_i \neq \mu_i$, and will be based on finite-dimensional problems.

5. p -SCLI- n for n -player Games

5.1. Definitions and examples

Let $\mathcal{Q}^{d_1, \dots, d_n}$ denote the set of n -player quadratic games, i.e. games comprised of n quadratic objectives $l_i : \mathbb{R}^{d_i} \rightarrow \mathbb{R}$, and $f_{\mathbf{A}, \mathbf{b}}(\mathbf{w}) \in \mathcal{Q}^{d_1, \dots, d_n}$ be a game with vector field $\mathbf{A}\mathbf{w} + \mathbf{b}$ as indicated in eq. 5. The following definition is a direct generalisation of the definition of p -SCLI algorithms given by Arjevani et al. (2016) to n -player games.

Definition 5 (p -SCLI- n optimisation algorithms for n -player games). *Let A be an optimisation algorithm for*

n -player quadratic games. Then \mathcal{A} is a p -stationary canonical linear iterative method for n -player games (p -SCLI- n) if there exist functions C_0, \dots, C_{p-1}, N from $\mathbb{R}^{d \times d}$ to $\mathbb{R}^{d \times d}$ -valued random variables, such that the following conditions are satisfied for all $f_{\mathbf{A}, \mathbf{b}}(\mathbf{w}) \in \mathcal{Q}^{d_1, \dots, d_n}$:

1. Given an initialisation $\mathbf{w}^0, \dots, \mathbf{w}^{p-1} \in \mathbb{R}^d$, the update rule at iteration $t \geq p$ is given by

$$\mathbf{w}^t = \sum_{i=0}^{p-1} C_i(\mathbf{A}) \mathbf{w}^{t-p+i} + N(\mathbf{A}) \mathbf{b} \quad (13)$$

2. $C_0(\mathbf{A}), \dots, C_{p-1}(\mathbf{A}), N(\mathbf{A})$ are independent from previous iterations
3. $\mathbb{E}C_i(\mathbf{A})$ are finite and simultaneously triangularisable

We will refer to the C_i as the coefficient matrices and N as the inversion matrix.

An important fact is that if $n = 1$, this definition becomes the same as the one given by Arjevani et al. (2016). A key difference, however, is that the Jacobian \mathbf{A} will generally not be symmetric for $n > 1$; only the blocks M_{ii} will be, and hence we may not assume the spectrum $\sigma(\mathbf{A})$ to be positive since it will generally be complex. Fortunately, several results from Arjevani et al. (2016) hold nevertheless, as discussed in Appendix C.1. Before introducing the results, let us give examples of algorithms used to optimise games that are p -SCLI- n , as evidenced by their update rule on quadratic games.

Simultaneous Gradient Descent (GD) The update rule is given by $\mathbf{w}_i^t = \mathbf{w}_i^{t-1} - \eta_i \nabla_{\mathbf{w}_i} l_i(\mathbf{w}^{t-1})$, which can be rewritten with $\boldsymbol{\eta} = \text{Diag}(\eta_1, \dots, \eta_n)$ as:

$$\begin{aligned} \mathbf{w}^t &= \mathbf{w}^{t-1} - \boldsymbol{\eta} (\mathbf{A} \mathbf{w}^{t-1} + \mathbf{b}) \\ &= (\mathbf{I} - \boldsymbol{\eta} \mathbf{A}) \mathbf{w}^{t-1} - \boldsymbol{\eta} \mathbf{b} \end{aligned} \quad (14)$$

This shows that simultaneous gradient descent is a 1-SCLI- n algorithm.

Simultaneous Momentum GD The update rule is $\mathbf{w}_i^t = \mathbf{w}_i^{t-1} - \eta_i \nabla_{\mathbf{w}_i} l_i(\mathbf{w}^{t-1}) + \beta_i (\mathbf{w}_i^{t-1} - \mathbf{w}_i^{t-2})$ which can be rewritten with $\boldsymbol{\beta} = \text{Diag}(\beta_1, \dots, \beta_n)$ and $\boldsymbol{\eta}$ as before:

$$\begin{aligned} \mathbf{w}^t &= \mathbf{w}^{t-1} - \boldsymbol{\eta} (\mathbf{A} \mathbf{w}^{t-1} + \mathbf{b}) + \boldsymbol{\beta} (\mathbf{w}^{t-1} - \mathbf{w}^{t-2}) \\ &= (\mathbf{I} - \boldsymbol{\eta} \mathbf{A} + \boldsymbol{\beta}) \mathbf{w}^{t-1} - \boldsymbol{\beta} \mathbf{w}^{t-2} - \boldsymbol{\eta} \mathbf{b} \end{aligned} \quad (15)$$

Therefore, simultaneous gradient descent with momentum is a 2-SCLI- n , if we assume $\boldsymbol{\beta}$ to be scalar (since we need the coefficient matrices $C_i(\mathbf{A})$ to be simultaneously triangularisable).

Extragradient (Korpelevich, 1976) The update rule is $\mathbf{w}_i^t = \mathbf{w}_i^{t-1} - \eta_i \nabla_{\mathbf{w}_i} l_i(\mathbf{w}^{t-1} - \boldsymbol{\eta} \mathbf{v}(\mathbf{w}^{t-1}))$, which can

be rewritten as:

$$\begin{aligned} \mathbf{w}^t &= \mathbf{w}^{t-1} - \boldsymbol{\eta} (\mathbf{A} (\mathbf{w}^{t-1} - \boldsymbol{\eta} (\mathbf{A} \mathbf{w}^{t-1} + \mathbf{b})) + \mathbf{b}) \\ &= (\mathbf{I} - \boldsymbol{\eta} \mathbf{A} + (\boldsymbol{\eta} \mathbf{A})^2) \mathbf{w}^{t-1} - (\mathbf{I} - \boldsymbol{\eta} \mathbf{A}) \boldsymbol{\eta} \mathbf{b} \end{aligned} \quad (16)$$

This shows that extragradient is a 1-SCLI- n .

Simultaneous Stochastic Gradient Descent The reasoning is the same as the one presented by Arjevani et al. (2016): we approximate $\nabla f_{\mathbf{A}, \mathbf{b}}(\mathbf{w}) = \mathbf{A} \mathbf{w} + \mathbf{b}$ with stochastic gradients $\mathbf{G}_\omega(\mathbf{w})$ and denote the error by $e_\omega(\mathbf{w}) = \mathbf{G}_\omega(\mathbf{w}) - (\mathbf{A} \mathbf{w} + \mathbf{b})$. Then the update rule for fixed $\boldsymbol{\eta}$ is given by

$$\begin{aligned} \mathbf{w}^t &= \mathbf{w}^{t-1} - \boldsymbol{\eta} \mathbf{G}_{\omega_{t-1}}(\mathbf{w}^{t-1}) \\ &= (\mathbf{I} - \boldsymbol{\eta} \mathbf{A}) \mathbf{w}^{t-1} - \boldsymbol{\eta} \mathbf{b} - \boldsymbol{\eta} e_{\omega_{t-1}}(\mathbf{w}^{t-1}) \end{aligned} \quad (17)$$

Under certain assumptions, e.g. if $e_\omega(\mathbf{w}) = \mathbf{A}_\omega \mathbf{w} + \mathbf{N}_\omega \mathbf{b}$ and $\mathbb{E} \mathbf{A}_\omega = \mathbb{E} \mathbf{N}_\omega = 0$, then the update rule becomes

$$\mathbf{w}^t = (\mathbf{I} - \boldsymbol{\eta} (\mathbf{A} + \mathbf{A}_{\omega_{t-1}})) \mathbf{w}^{t-1} - \boldsymbol{\eta} (\mathbf{I} + \mathbf{N}_{\omega_{t-1}}) \mathbf{b} \quad (18)$$

and we get a 1-SCLI- n .

One last definition is required before we introduce the p -SCLI lower bounds. Our definition generalises that of Arjevani et al. (2016).

Definition 6 (Consistency of p -SCLI- n optimisation algorithms). Let $\mathcal{Q}_{\mathbf{A}}^{d_1, \dots, d_n} \subseteq \mathcal{Q}^{d_1, \dots, d_n}$ denote the set of quadratic n -player games with non-singular Jacobian \mathbf{A} (see eq. 5). Then \mathcal{A} is consistent with respect to \mathbf{A} if for any game $f_{\mathbf{A}, \mathbf{b}} \in \mathcal{Q}_{\mathbf{A}}^{d_1, \dots, d_n}$ and any initialisation, \mathcal{A} converges to a stationary point of $f_{\mathbf{A}, \mathbf{b}}$ or equivalently if the sequence of iterates (\mathbf{w}^t) (see eq. 13) satisfies

$$\mathbf{w}^t \rightarrow -\mathbf{A}^{-1} \mathbf{b} \quad (19)$$

Equivalently, as Arjevani et al. (2016) argue in their section 3.1, consistency with respect to some invertible Jacobian \mathbf{A} is equivalent to having \mathcal{A} converge on $f_{\mathbf{A}, \mathbf{b}}$ and

$$\sum_{i=0}^{p-1} \mathbb{E} C_i(\mathbf{A}) = \mathbf{I}_d + \mathbb{E} N(\mathbf{A}) \mathbf{A} \quad (20)$$

Note that all three examples of optimisation algorithms discussed in this subsection satisfy eq. 20.

5.2. Parametric lower bound for p -SCLI- n with scalar inversion matrix

We are now ready to introduce the lower bound for p -SCLI- n methods with scalar inversion matrix.

Proposition 7. Let \mathcal{A} be a p -SCLI- n algorithm with scalar inversion matrix for optimising games over $\mathbb{R}^{d_1} \times \dots \times \mathbb{R}^{d_n}$.

Then for quadratics $f_{\mathbf{A},\mathbf{b}} \in \mathcal{Q}^{d_1, \dots, d_n}$, if \mathbf{A} is consistent with respect to \mathbf{A} and if $0 \notin \sigma(\mathbf{A})$, we have the following lower bound on the (linear) rate of convergence ρ :

$$\rho \geq \frac{\sqrt[p]{\kappa} - 1}{\sqrt[p]{\kappa} + 1} = 1 - \frac{2}{\sqrt[p]{\kappa} + 1} \quad (21)$$

where the condition number κ is defined as $\kappa \triangleq \frac{\max |\sigma(\mathbf{A})|}{\min |\sigma(\mathbf{A})|}$ where $\sigma(\mathbf{A})$ is the spectrum of \mathbf{A} .

While this lower bound is valid for p -SCLI- n methods on any single quadratic game where the assumptions apply, it also gives us a more general result for problem classes containing quadratics as an immediate corollary, similarly to the lower bounds of the previous section.

Theorem 8. *For any problem class containing quadratic games, there exists a game such that the lower bound given in eq. 21 holds for a p -SCLI- n method.*

Corollary 9 (Iteration complexity bound). *For the same problem classes and under the same assumptions as Theorem 8, the minimal number of steps t to satisfy $\max_{i=0, \dots, p-1} \|\mathbb{E}\mathbf{w}^{t+i} - \mathbb{E}\mathbf{w}^*\| < \epsilon$ is given by*

$$t \geq \left(\frac{\sqrt[p]{\kappa} - 1}{2} \right) \log \left(\frac{C}{\epsilon} \right) \quad (22)$$

for some strictly positive constant C .

See Appendix C.1 for the proof of Prop. 7, and Appendix D for the iteration complexity bound's proof. Interestingly, by setting $n = 1$, this bound captures the 1-player case for μ -strongly convex objectives with L -Lipschitz gradients, where $\kappa = \frac{\max \sigma(\mathbf{A})}{\min \sigma(\mathbf{A})} = L/\mu$, verifying the intuition discussed at the end of the previous section, and showing that this technique yields a tight lower bound for $n = 1$. Moreover, this form is valid for n -player games (and min-max problems) in finite dimension, and κ arises naturally from the spectral properties of the update rules of the p -SCLI- n methods and is lower bounded by 1. Additionally, our bounds are valid for some stochastic methods. In single-objective optimisation (i.e. $n = 1$), while linear rates are not achievable for general stochastic problems, for which the worst-case bounds are sublinear, under certain conditions linear rates are possible (Loizou and Richtárik, 2017). Conditions of this type can be satisfied in over-parameterised neural networks (Vaswani et al., 2019). Hence, our linear lower bounds may be useful even for stochastic problems.

However, while the moduli in the $n > 1$ case allow us to handle complex spectra and matches the classical definition of condition number from linear algebra, several analyses have shown that not only the modulus, but also the relative size of the real and imaginary parts of elements of the spectrum matter (Mescheder et al., 2017; Gidel et al., 2019b). Such an analysis may yield more expressive bounds, but

is out of the scope of this work. We will nevertheless give a more explicit form of the bound for 2-player games for which $d_1 = d_2$ that will make the μ_i and L_i appear.

Some explicit bounds for p -SCLI-2 with $d_1 = d_2$ Prop. 7 may be used to derive lower bounds for 2-player games for which $d_1 = d_2$. These bounds depend on the value of the μ_i and L_i defined as in eq. 6. Namely, let

$$\begin{aligned} \Delta_\mu &= (\mu_1 + \mu_2)^2 - 4(\mu_1\mu_2 + \mu_{12}^2) \\ &= (\mu_1 - \mu_2)^2 - 4\mu_{12}^2 \end{aligned} \quad (23)$$

$$\begin{aligned} \Delta_L &= (L_1 + L_2)^2 - 4(L_1L_2 + L_{12}^2) \\ &= (L_1 - L_2)^2 - 4L_{12}^2 \end{aligned} \quad (24)$$

Table 1 gives lower bounds on the condition number that may then be plugged into eq. 21 to get lower bounds on two-players games corresponding to min-max problems (and are therefore lower bounds for general 2-player games).

Table 1. Lower Bounds on the Condition Number

		$\Delta_\mu < 0$	$\Delta_\mu \geq 0$
$\Delta_L < 0$	$\kappa = \sqrt{\frac{L_1L_2 + L_{12}^2}{\mu_1\mu_2 + \mu_{12}^2}}$	$\kappa \geq 2 \sqrt{\frac{L_1L_2 + L_{12}^2}{\mu_1 + \mu_2 - \sqrt{\Delta_\mu}}}$	
$\Delta_L \geq 0$	$\kappa \geq \frac{1}{2} \frac{L_1 + L_2 + \sqrt{\Delta_L}}{\sqrt{\mu_1\mu_2 + \mu_{12}^2}}$		$\kappa \geq \frac{L_1 + L_2 + \sqrt{\Delta_L}}{\mu_1 + \mu_2 - \sqrt{\Delta_\mu}}$

See Appendix C.2 for the counterexample in \mathcal{P} leading to these bounds. This result resolves the issues raised in the discussion of the domino bounds as it uses all of the μ_i and L_i , and is proven from finite-dimensional problems. In fact, if one sets $\mu_1 = L_1$ and $\mu_2 = L_2$ such that μ_1 and μ_2 are small (in particular, smaller than μ_{12}) and L_{12} is large (in particular, larger than L_1, L_2), table 1 yields $\kappa = \sqrt{\frac{\mu_1\mu_2 + L_{12}^2}{\mu_1\mu_2 + \mu_{12}^2}}$, which coincides with the κ from Thm. 3. More generally, for $p = 1$, if both Δ_L and Δ_μ are negative, we get a tighter bound from the p -SCLI-2 formalism for 1-SCLI-2 methods that also satisfy the two-step linear span assumption than from Thm. 3. Finally, it provides the same plug-and-play convenience as p -SCLI to derive bounds for a large class of algorithms that may not satisfy the first-order black box assumption. On the other hand, the bounds may not be tight for $p \geq 3$, as it was the case for single-objective p -SCLI (Arjevani et al., 2016).

An interesting case is $p = 2$: for 2-SCLI-2 methods that satisfy the two-step linear span assumption such as negative momentum, the rate stemming from the p -SCLI-2 analysis appears to be smaller than the rate from the improved domino bound. This may be because the proof techniques used in our generalisation of p -SCLI yield bounds that can be improved for $p > 1$. In particular, a key difference between Thm. 3 and 8 is that as explained in the background section, lower bounds generally need not apply to every

single problem within a class. However, due to the proof technique used, the bound given in Prop. 7 has to apply to every single quadratic game where the assumptions hold. Because some games might be inherently easier than others, such games may bottleneck the bound in Prop. 7 and Thm. 8, and introduce looseness. Indeed, first, the number of players does not appear in Prop. 7’s proof, meaning that the proof yielded a bound that should hold for $n = 1$ which might inherently have faster convergence than $n > 1$ due to having to optimise only one objective over one set of parameters. Second, in the proof, because we do not impose a structure for the quadratic games, the bound has to hold for purely cooperative games (that is, the objective of each player does not depend on other players’ parameters), in which case the situation is analogous to n single-objective optimisation problems, where we may expect the convergence to be faster than in games with interactions between players. As such, it is worth considering whether fixing the number of players and the structure of the games earlier in the proof could improve the bound.

6. Conclusion

In this work, we provide linear lower bounds and condition numbers for any problem class containing quadratic or bilinear games. We give a lower bound on the rate of convergence of first-order black box methods for 2-player games (which directly applies to min-max and saddle point problems) satisfying the two-step linear span assumption by generalising Nesterov’s lower bound for the optimisation of strongly convex, smooth convex objectives to 2-player games (R.Q. 1) and constructing a novel class of hard problems using spectral properties of a class of operators in Banach algebras. Moreover, we generalise the framework of p -SCLI, which requires symmetricity of the Hessian in single-objective optimisation, to provide a bound for a large class of optimisers for n -player games by extending the results of p -SCLI to quadratic games with non-symmetric Jacobian, which for $n = 1$ recovers (Arjevani et al., 2016)’s tight bounds. We then give explicit bounds for 2-player games, which apply to min-max and saddle point problems (R.Q. 2). Finally, we derived formulations for the condition number that matched (in the case of the first domino bound), or were more general (in the case of the improved domino bound, and p -SCLI- n and p -SCLI-2 bounds) than the existing ones in the upper bound literature (R.Q. 3). As in the single-objective case, our bounds and condition numbers suggest that optimisers may converge faster on games for which the eigenvalues are at a similar, remote distance from the origin (e.g. on a circle) than on games for which some eigenvalues are close to and others are far from 0.

Following the initial release of this work, several other authors have built upon the bounds presented in this paper.

The bound from Theorem 3 is tight on the class of smooth strongly-convex-strongly-concave games, as it is matched by the upper bound presented by Fallah et al. (2020), and up to logarithmic factors, by the upper bound of Lin et al. (2020). Additionally, this same lower bound was also shown to be tight on the class of bilinear games with non-singular Jacobian, which are merely convex-concave, by Prop. 4 of Azizian et al. (2020). As a corollary, our results establish the optimality of those methods on the aforementioned classes of problems.

However, several directions remain to be explored. For example, we raised the question of whether the p -SCLI- n bound could be tightened for $p > 1$ by improving the proof technique, especially given that we are not aware of faster rates of convergence in the literature than those of $p = 1$ or those of the improved domino bound. Moreover, we would like to present a more exhaustive overview of the extension of p -SCLI, and discuss the resulting upper and lower bounds for various commonly used algorithms. In particular, we would like to extend our lower bounds to p -SCLI- n with diagonal inversion matrices, as Arjevani et al. (2016) did in the p -SCLI framework, and provide bounds in the 2-player case when $d_1 \neq d_2$. Furthermore, we believe tighter bounds may be derived, for example by adding constraints on the C_i or by looking not only at the modulus of the eigenvalues but also at their arguments, as done by Gidel et al. (2019b), since we know that the relative size of the imaginary part and real part (even at fixed modulus) affects the dynamics in games (Mescheder et al., 2017). Finally, it would be important to understand when and how linear convergence is possible in non strongly-convex-strongly-concave settings. We plan on exploring several of these directions in future work.

Acknowledgements

The authors would like to thank Damien Scieur for useful discussions and feedback. This work was partially supported by the FRQNT new researcher program (2019-NC-257943), the NSERC Discovery grant (RGPIN-2019-06512), a startup grant by IVADO, a Microsoft Research collaborative grant and a Canada CIFAR AI chair. Gauthier Gidel was partially supported by a Borealis AI Graduate Fellowship.

REFERENCES

- Yossi Arjevani, Shai Shalev-Shwartz, and Ohad Shamir. On lower and upper bounds in smooth and strongly convex optimization. *The Journal of Machine Learning Research*, 17(1):4303–4353, 2016.
- Waïss Azizian, Damien Scieur, Ioannis Mitliagkas, Simon Lacoste-Julien, and Gauthier Gidel. Accelerating smooth games by manipulating spectral shapes. *arXiv preprint*

- arXiv:2001.00602*, 2020.
- David Balduzzi, Sebastien Racaniere, James Martens, Jakob Foerster, Karl Tuyls, and Thore Graepel. The mechanics of n-player differentiable games. In *International Conference on Machine Learning*, pages 363–372, 2018.
- Michele Benzi and Gene H Golub. A preconditioner for generalized saddle point problems. *SIAM Journal on Matrix Analysis and Applications*, 26(1):20–41, 2004.
- Gilles Brassard and Paul Bratley. *Fundamentals of algorithms*. 1996.
- Sébastien Bubeck et al. Convex optimization: Algorithms and complexity. *Foundations and Trends® in Machine Learning*, 8(3-4):231–357, 2015.
- Antonin Chambolle and Thomas Pock. A first-order primal-dual algorithm for convex problems with applications to imaging. *Journal of mathematical imaging and vision*, 40(1):120–145, 2011.
- George HG Chen and R Tyrrell Rockafellar. Convergence rates in forward–backward splitting. *SIAM Journal on Optimization*, 7(2):421–444, 1997.
- Yunmei Chen, Guanghai Lan, and Yuyuan Ouyang. Accelerated schemes for a class of variational inequalities. *Mathematical Programming*, 165(1):113–149, 2017.
- Constantinos Daskalakis, Andrew Ilyas, Vasilis Syrgkanis, and Haoyang Zeng. Training GANs with optimism. In *International Conference on Learning Representations*, 2018.
- Ronald G Douglas. *Banach algebra techniques in operator theory*, volume 179. Springer Science & Business Media, 2012.
- Simon S Du, Jianshu Chen, Lihong Li, Lin Xiao, and Dengyong Zhou. Stochastic variance reduction methods for policy evaluation. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1049–1058. JMLR. org, 2017.
- Alireza Fallah, Asuman Ozdaglar, and Sarath Pattathil. An optimal multistage stochastic gradient method for minimax problems. *arXiv preprint arXiv:2002.05683*, 2020.
- Gauthier Gidel, Hugo Berard, Gaëtan Vignoud, Pascal Vincent, and Simon Lacoste-Julien. A variational inequality perspective on generative adversarial networks. In *International Conference on Learning Representations*, 2019a.
- Gauthier Gidel, Reyhane Askari Hemmat, Mohammad Pezeshki, Rémi Le Priol, Gabriel Huang, Simon Lacoste-Julien, and Ioannis Mitliagkas. Negative momentum for improved game dynamics. In *The 22nd International Conference on Artificial Intelligence and Statistics, AISTATS 2019*, pages 1802–1811, 2019b.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.
- Patrick T Harker and Jong-Shi Pang. Finite-dimensional variational inequality and nonlinear complementarity problems: a survey of theory, algorithms and applications. *Mathematical programming*, 48(1-3):161–220, 1990.
- Seung-Jean Kim and Stephen Boyd. A minimax theorem with applications to machine learning, signal processing, and finance. *SIAM Journal on Optimization*, 19(3):1344–1367, 2008.
- GM Korpelevich. The extragradient method for finding saddle points and other problems. *Matecon*, 12:747–756, 1976.
- Tianyi Lin, Chi Jin, Michael Jordan, et al. Near-optimal algorithms for minimax optimization. *arXiv preprint arXiv:2002.02417*, 2020.
- Nicolas Loizou and Peter Richtárik. Momentum and stochastic momentum for stochastic gradient, newton, proximal point and subspace descent methods. *arXiv preprint arXiv:1712.09677*, 2017.
- Lars Mescheder, Sebastian Nowozin, and Andreas Geiger. The numerics of GANs. In *Advances in Neural Information Processing Systems*, pages 1825–1835, 2017.
- Arkadii Semenovich Nemirovsky and David Borisovich Yudin. Problem complexity and method efficiency in optimization. 1983.
- AS Nemirovsky. Information-based complexity of linear operator equations. *Journal of Complexity*, 8(2):153–175, 1992.
- Yurii Nesterov. *Introductory lectures on convex optimization: A basic course*, volume 87. Springer Science & Business Media, 2004.
- Yurii E Nesterov. A method for solving the convex programming problem with convergence rate $O(1/k^2)$. In *Dokl. akad. nauk Sssr*, volume 269, pages 543–547, 1983.
- Yuyuan Ouyang and Yangyang Xu. Lower complexity bounds of first-order methods for convex-concave bilinear saddle-point problems. *arXiv preprint arXiv:1808.02901*, 2018.

Balamurugan Palaniappan and Francis Bach. Stochastic variance reduction methods for saddle-point problems. In *Advances in Neural Information Processing Systems*, pages 1416–1424, 2016.

Lewis Fry Richardson. IX. the approximate arithmetical solution by finite differences of physical problems involving differential equations, with an application to the stresses in a masonry dam. *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character*, 210(459-470): 307–357, 1911.

Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training GANs. In *Advances in neural information processing systems*, pages 2234–2242, 2016.

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, AISTATS 2019*, pages 1195–1204, 2019.

Hongxing Wang, Chaoqun Weng, and Junsong Yuan. Multi-feature spectral clustering with minimax optimization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4106–4113, 2014.

Fuzhen Zhang, editor. *The Schur Complement and Its Applications*, volume 4 of *Numerical Methods and Algorithms*. Springer-Verlag, New York, 2005. ISBN 978-0-387-24271-2.