



HAL
open science

Global Min-Max Computation for α -Hölder Zero-Sum Games

Aurélien Delage, Olivier Buffet, Jilles Dibangoye

► **To cite this version:**

Aurélien Delage, Olivier Buffet, Jilles Dibangoye. Global Min-Max Computation for α -Hölder Zero-Sum Games. GAIW 2023 - 5th Games, Agents, and Incentives Workshop, May 2023, Londres (London), United Kingdom. pp.1-9. hal-04382817

HAL Id: hal-04382817

<https://inria.hal.science/hal-04382817>

Submitted on 9 Jan 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Global Min-Max Computation for α -Hölder zero-sum Games

Aurélien Delage
Univ Lyon, INSA Lyon, Inria, CITI
69621 Villeurbanne, France
aurelien.delage@insa-lyon.fr

Olivier Buffet
Univ. Lorraine, INRIA, CNRS, LORIA
F-54000, Nancy, France
olivier.buffet@inria.fr

Jilles-Steeve Dibangoye
Univ Lyon, INSA Lyon, Inria, CITI
69621 Villeurbanne, France
jilles-steeve.dibangoye@insa-lyon.fr

ABSTRACT

Min-max optimization problems recently arose in various settings. From Generative Adversarial Networks (GANs) to aerodynamic optimization through Game Theory, the assumptions on the objective function vary. Motivated by the applications to deep learning and especially GANs, most recent works assume differentiability to design local search algorithms such as Gradient Descent Ascent (GDA). In contrast, this work will only require α -Hölder properties to tackle general game-theoretic problems with poor continuity assumptions. Focusing on the example of problems in which max and min optimization variables live in simplices, we provide a simple algorithm, based on Deterministic Optimistic Optimization (DOO), relying on an outer min-optimization using the solutions of an inner max-optimization. The algorithm is shown to converge in finite time to an ϵ -global optimum. Experimental validations are given and the time complexity of our algorithm is studied.

ACM Reference Format:

Aurélien Delage, Olivier Buffet, and Jilles-Steeve Dibangoye. 2023. Global Min-Max Computation for α -Hölder zero-sum Games. In *Appears at the 5th Games, Agents, and Incentives Workshop (GAIW 2023). Held as part of the Workshops at the 21st International Conference on Autonomous Agents and Multiagent Systems., London, England, May-June 2023, IFAAMAS*, 9 pages.

1 INTRODUCTION

There is a growing interest in min - max problems from various communities. Generative Adversarial Networks [22, 25], fair statistical inference [14, 26], robust decision-making [5], and general resource allocation [8] witness its importance for machine learning. While it has been known for a long time that formulating real-life problems as games is relevant to economics [11], physical phenomena [4] are also interestingly linked to such an optimization problem. Basically, any situation in which two entities are trying to optimize different criteria by interacting with a common system might be seen through a game-theoretic perspective. The players' actions having heterogeneous influences on the system, min - max optimization is a problem whose hypotheses and, thus, difficulty substantially vary depending on the context.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

Appears at the 5th Games, Agents, and Incentives Workshop (GAIW 2023). Held as part of the Workshops at the 21st International Conference on Autonomous Agents and Multiagent Systems., Abramowitz, Aziz, Ceppi, Dickerson, Hosseini, Lev, Mattei, Tsang, Wqs, Zick (Chairs), May-June 2023, London, England. © 2023 Copyright held by the owner/author(s). . . \$ACM ISBN 978-x-xxxx-xxxx-x/YY/MM

As emphasized by Daskalakis [6], applications of game theory to Deep Learning require modeling more and more complex interactions, making the usual convex-concavity or even differentiability assumptions unfortunately too restrictive. This paper thus provides a first answer to tackle such problems, only requiring mild α -Hölder properties.

Moving on to a formal problem definition, for any given function $f : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$ (where $\mathcal{X} \subset \mathbb{R}^n$ and $\mathcal{Y} \subset \mathbb{R}^p$), the general min - max problem is to compute at least one pair $(\mathbf{x}^*, \mathbf{y}^*)$ such that

$$\mathbf{x}^* \in \arg \min_x \max_y f(\mathbf{x}, \mathbf{y}), \quad (1)$$

$$\mathbf{y}^* \in \arg \max_y \min_x f(\mathbf{x}, \mathbf{y}). \quad (2)$$

The simplest case, introduced by von Neumann [29], assumes that (i) \mathcal{X} and \mathcal{Y} are respectively the unit simplices of \mathbb{R}^n and \mathbb{R}^p , and (ii) f is bilinear. Then, min - max equals max - min, the problem corresponds to finding Nash Equilibrium of a 2-player 0-sum Normal-Form Game, and is solved by a simple linear program in polynomial time [27]. Even though Sion [28] showed that (i) and (ii) can be respectively replaced by (iii) \mathcal{X} and \mathcal{Y} convex and (iv) f convex-concave with min - max still being equal to max - min, such equality and simplicity of resolution does not hold in the general case.

Relaxing assumptions (iii) and (iv) on f leads to various concerns. If f is not convex or not concave, Nash Equilibria do not necessarily exist [6]. Even if they do, the computational cost of finding one might be very high, as deciding whether a Nash Equilibrium exists when the game is not convex-concave is NP-Hard in general [7].

Assuming only differentiability (but not that a Nash equilibrium necessarily exists), the most common approach to solve the min - max problem is to alternate between gradient ascents and descents to respectively comply with max and min's will. The resulting algorithm is known as *Gradient Descent Ascent (GDA)*. However, this can fail [13], even for some simple zero-sum bilinear games [17, 18].

For this reason, under continuous differentiability assumption, a First-order Nash Equilibrium (FNE) solution concept was defined through the first order Taylor expansion approximation of the function, and the existence of at least one FNE is guaranteed for twice differentiable functions [21]. Recent work introduced modifications to the GDA algorithm [13, 21, 25] to provide algorithms asymptotically converging towards a FNE in a number of gradient evaluations going from $\mathcal{O}(\epsilon^{-2})$ to $\mathcal{O}(\epsilon^{-6})$ [24] through $\mathcal{O}(\epsilon^{-4})$ [15, 16], depending on the assumptions made on f . As a matter of fact, computing ϵ -FNEs has been shown to be in ExpTime with respect to either $1/\epsilon$, the smoothness of the game, or its dimensionality [7].

The intensive study of gradient-based approaches might be due to its utility for deep learning. Thus, min - max problems without,

at least, differentiability assumptions were not studied even though they naturally appear [6, Figure 1,b][1, 2]. Since we are interested in games with poor continuity properties, we will focus on the min - max solution concept that is well-defined and exists.

Encouragingly, the min - max problem always admits a solution (by f 's continuity) and the optima of a “local”, *i.e.* players can only deviate within a ball of radius δ from their strategy, min - max version of the problem [15] are strongly linked to the stable limit point of GDA [15]. Besides, the min - max computation makes sense as it corresponds to finding *security values* for player min by searching for the most rewarding strategy that she can announce to player max and the resulting problem is called a *zero-sum Stackelberg game*.

For now, relaxing assumptions faces the instability and divergent behavior of gradient-based approaches. To tackle this problem, we build on a global optimizing algorithm for functions with light continuity assumptions, namely Munos' Deterministic Optimistic Optimization (DOO) [19], to ϵ -optimally solve the min - max problem in finite time for functions with only mild α -Hölder continuity properties. This allows to deal with a larger class of games, starting with the ones including entropy measures in the payoff function [6, Figure 1,b][1, 2].

After providing some technical background regarding game theory and DOO in Section 2, Section 3 introduces two adaptations of DOO in order to provide a finite-time convergent algorithm whenever the variables live in simplices. Section 4 then derives an algorithm to ϵ -optimally solve the min-max problem, and this resulting algorithm is experimentally studied in Section 5.

2 BACKGROUND

Note : Vectors \mathbf{x} of \mathbb{R}^n are written in bold, and the k -th component of \mathbf{x} is x^k .

First, we start by providing some necessary background about game theory and insights about the relevance of the α -Hölder condition in games.

2.1 Games and solution concepts

Definition 2.1 (Two player zero-sum Game). A two player zero-sum game is defined by a tuple $\langle \mathcal{A}^1, \mathcal{A}^2, f \rangle$ where, for each $i \in \{1, 2\}$, there exists $p_i \in \mathbb{N}$ such that \mathcal{A}^i is a compact subset of \mathbb{R}^{p_i} , and $f : \mathcal{A}^1 \times \mathcal{A}^2 \rightarrow \mathbb{R}$ is continuous but not necessarily differentiable. The game will be called a *normal-form game* if (i) \mathcal{A}^1 and \mathcal{A}^2 are respectively the unit simplex S_{p_1} of \mathbb{R}^{p_1} and the unit simplex S_{p_2} of \mathbb{R}^{p_2} and (ii) f is a matrix belonging to \mathcal{M}_{p_1, p_2} .

A simple, yet interesting, example consists in a zero-sum normal-form game $\langle S_2, S_2, M \rangle$ in which player 1 (resp. 2) also wants to minimize (resp. maximize) the entropy of her mixed strategy, while not degrading too much their expected payoff.

Example 2.2 ([6]). Consider the classical matching pennies game given by the payoff matrix M

$$M = \begin{pmatrix} 1/2 & -1/2 \\ -1/2 & 1/2 \end{pmatrix}, \quad (3)$$

and the bilinear payoff function $\mathbf{x}^\top \cdot M \cdot \mathbf{y}$. The payoff function of the modified matching pennies game in which player 1 is rewarded

for a high entropy strategy but 2 is penalized for a high entropy strategy is formally given by $f(\mathbf{x}, \mathbf{y}) = \mathbf{x}^\top \cdot M \cdot \mathbf{y} + x_1 \log(x_1) + x_2 \log(x_2) + y_1 \log(y_1) + y_2 \log(y_2)$. In this zero-sum game, no Nash equilibrium exists, and the payoff function is neither convex nor concave. Still, the payoff function is α -Hölder for any $\alpha \in]0, 1[$.

Definition 2.3 (α -Hölder condition). For any norm $\|\cdot\|$, a function $f : \mathcal{X} \rightarrow \mathbb{R}$ satisfies the α -Hölder condition with respect to $\|\cdot\|$ if and only if there exists $\alpha \in]0, 1[$ and $C \in \mathbb{R}^+$ such that

$$\forall (\mathbf{x}, \tilde{\mathbf{x}}) \in \mathcal{X}^2, |f(\mathbf{x}) - f(\tilde{\mathbf{x}})| \leq C \|\mathbf{x} - \tilde{\mathbf{x}}\|^\alpha. \quad (4)$$

As mentioned in the introduction, Nash Equilibrium points do not necessarily exist, in which case max - min $>$ min - max, even in a local sense [15]. Players thus can not agree on a joint strategy. Thus, they might want to search for the individual strategy that has the greatest value against any strategy of their opponent, *i.e.* (for player 1) compute $\arg \min_{\mathbf{x}} \max_{\mathbf{y}} f(\mathbf{x}, \mathbf{y})$.

Since this paper aims to tackle games with poor continuity properties of the payoff function, we will only assume that f is α -Hölder¹.

2.2 Deterministic Optimimistic Optimization (DOO)

Let $f : \mathcal{X} \rightarrow \mathbb{R}$ be a λ -Lipschitz function with respect to any semi-metric l (*i.e.* $\forall (\mathbf{x}, \mathbf{x}') \in \mathcal{X}^2, |f(\mathbf{x}) - f(\mathbf{x}')| \leq \lambda l(\mathbf{x}, \mathbf{x}')$) defined over $\mathcal{X} \subset \mathbb{R}^n$.

If \mathcal{X} is bounded, DOO ϵ -optimally solves $\min_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})$ for a given error $\epsilon > 0$ by creating a non-uniform covering of the domain with a finite number of cells where the variations of the function are controlled.

Given any semi-metric l , a cell S and point $\tilde{\mathbf{x}}$ in S , f 's Lipschitz continuity with respect to l allows optimistically bounding (*i.e.* lower bounding) its value within S by $f(\tilde{\mathbf{x}}) - \lambda \text{Diam}(S)$, where $\text{Diam}(S) = \max_{\mathbf{x}, \mathbf{y} \in S} l(\mathbf{x}, \mathbf{y})$. The smaller the diameter of the cell, the closer the optimistic bound is to the values of f in it.

Let us assume that \mathcal{X} is such that (i) there is an analytic way to create a first covering with cells, and (ii) each cell can again be covered with children-cells. The algorithm (given in Algorithm 1) starts with the first covering of \mathcal{X} (line 2), and computes the optimistic bound in every cell (line 3). Then, each iteration consists in selecting the most promising cell according to the optimistic bound (line 6), covering it (line 7), and computing the optimistic bound of each new cell (line 8). The final returned value is the encountered point with the lowest value. Figure 1 gives an illustration of an execution of DOO.

If the cells are “well-formed” (following Assumption 1) and their diameter decreases w.r.t their depth in the tree (following Assumption 2), the approximation error shrinks with the radius around the most promising area as the algorithm iterates. Thus, the difference $r(n)$ between the highest value of f and DOO's returned value after performing n coverings of cells with children cells can be bounded [19].

ASSUMPTION 1 (VALID CELL (ASSUMPTION 4 IN [19])). *There exists $v \in \mathbb{R}^{+,*}$ such that for any cell $S \subset \mathcal{X}$, there exists $\mathbf{x} \in S$ such that*

¹for the sake of simplicity, we only expose the results for α -Hölder functions even though [19] presents the result for Lipschitz functions with respect to any semi metric l (not just $\|\cdot\|^\alpha$).

Algorithm 1: DOO

```

1 Fct DOO( $\mathcal{X} \rightarrow \mathbb{R}; x \mapsto f(x)$ ),  $n, l$ )
   input :  $f : \mathbb{R} \rightarrow \mathbb{R}$  is a  $\alpha$ -Hölder function with respect
           to a semi-norm  $l$ 
   input :  $C$  a  $\alpha$ -Hölder constant of  $f$ 
2 Initialize  $\mathcal{I}$  and  $(S_i)_{i \in \mathcal{I}}$  s.t.  $\mathcal{X} \subseteq \cup_{i \in \mathcal{I}} S_i$ 
   /* The diameter of a set  $S_i$  is noted  $\rho_i$ . */
3  $\forall i \in \mathcal{I}$ , compute  $f(x_i) - C\rho_i^\alpha$ 
4  $time \leftarrow 0$ 
5 while  $time \leq n$  do
6    $i^* \leftarrow \arg \min_{i \in \mathcal{I}} f(x_i) - C\rho_i^\alpha$ 
7   Recover  $S_{i^*}$  by  $\cup_{j \in \mathcal{I}^*} S_j$  ( $\supseteq S_{i^*}$ )
    $\forall j \in \mathcal{I}^*, x_j \leftarrow Repr(S_j)$ 
   /* Here,  $Repr(S_j)$  is any function returning
       a point that belongs to  $S_j$  and to  $\mathcal{X}$ . */
8    $\mathcal{I} \leftarrow [\mathcal{I} \setminus \{i^*\}] \cup \mathcal{I}^*$ 
9    $time \leftarrow time + 1$ 
10 return  $\langle \arg \& \min_{x_i} f(x_i) \rangle$ 

```

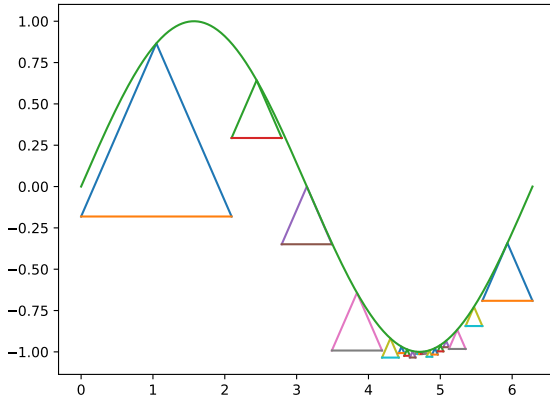


Figure 1: Representation of DOO's execution to minimize $x \mapsto \sin(x)$ on $[0, 2\pi]$. The interval is tiled with cells (i.e. here, intervals) of different sizes where the function \sin is lower-bounded by "Lipschitz cones", themselves lower-bounded by a constant. The cones, along with their constant lower bounds, form the triangular shapes.

$\mathcal{B}_l(x, vDiam(S)) \subset S$ where $\mathcal{B}_l(x, \rho) = \{y \in X | l(x, y) \leq \rho\}$ denote the ball of radius ρ centered in x for the semi-metric l .

ASSUMPTION 2. There exists a decreasing sequence $\delta(h)$ such that for any depth $h \geq 0$, for any cell $S_{h,i}$ of depth h in the tree, $Diam(S_{h,i}) \leq \delta(h)$.

Remark 2.4. Fortunately, for any given norm $\|\cdot\|$, one can remark that (i) an α -Hölder function with constant C is exactly a C -Lipschitz function with respect to the semi-metric $\|\cdot\|^\alpha$, and (ii) balls $\mathcal{B}_{\|\cdot\|^\alpha}(x, \rho)$ are valid cells. We might thus be able to adapt

Munos' work to tackle min - max problems for α -Hölder functions, using simple balls $\mathcal{B}_{\|\cdot\|^\alpha}(\cdot, \cdot)$.

The next two sections respectively show (i) how DOO can be applied to the special case of optimization over a simplex and (ii) how DOO can be adapted to solve a min - max problem.

3 FINITE-TIME CONVERGENT DOO FOR SIMPLEX SPACES

In this section, we come back to vanilla minimization (rather than min - max optimization) and provide two modifications to the original DOO algorithm in order to match usual requirements of game theory. First, we require our algorithm to converge in finite time to an ϵ -optimum instead of bounding the regret for a given time budget. This leads to introducing a pessimistic bound which will also be useful later as a pruning criterion. Secondly, we show that DOO can be used when the search space is a simplex by applying iterative simplex discretization methods.

3.1 Modifying the stopping criterion

The initial analysis of DOO's complexity bounds its error with respect to a given budget n . We change viewpoint to show that, for any $\epsilon > 0$, one can derive a stopping criterion that is reached in finite time. To do so, we introduce upper and lower bounds of the optimal value, which will monotonically shrink as the algorithm iterates. The algorithm returns whenever the difference between the maximum value of the upper-bound and the maximum value of the lower-bound is smaller than ϵ .

Definition 3.1 (Upper and lower bounds). Let f be an α -Hölder function and C be an α -Hölder constant of f . Let $(\mathcal{B}_{\|\cdot\|^\alpha}(x_i, \rho_i))_i$ be a set of cells partitioning f 's domain. Then, $\min_{x \in \mathcal{X}} f(x)$ is upper- and lower-bounded respectively by $\min_{x_i} f(x_i)$ and $\min_{x_i} f(x_i) - C\rho_i^\alpha$.

A straightforward adaptation of Munos' regret-bounding proof leads to the following convergence theorem.

THEOREM 3.2 (ADAPTED FROM [20]). Let $\epsilon > 0$ be a given error. Let $(\mathcal{B}_{\|\cdot\|^\alpha}(x_i, \rho_i))_i$ be the set of cells covering *[ad] partitioning->covering* f 's domain, updated as the algorithm iterates. The quantity $|\min_i [f(x_i) - C\rho_i^\alpha] - \min_i [f(x_i)]|$ decreases throughout the optimization process and is smaller than ϵ after a finite number of iterations.

Remark 3.3. One could derive a precise complexity bound, but at the cost of introducing function-dependent topological constants whose computation is, in general, even harder than the optimization problem. Instead, Theorem 4.3 (Section 4.1) will provide an upper bound of the complexity, based on the worst-case scenario of optimizing a constant function.

Note that (i) $\min_i f(x_i) - C\rho_i^\alpha$ and $\min_i f(x_i)$ have no reason to be attained in the same cell; and (ii) Theorem 3.2 leads to simply modifying line 5 of Algorithm 1 to stop whenever the gap between bounds is less than ϵ . The parameter n in Algorithm 1 is thus replaced by ϵ .

3.2 DOO for simplex spaces

Subdivision process. We now detail a process to subdivide the n -dimensional unit simplex S_n using smaller and smaller hypercubes as the process iterates. Simplex decomposition techniques are not new since Edelsbrunner and Grayson [9], and Paulavičius and Žilinskas [23] already introduced some and Paulavičius and Žilinskas did so to perform Lipschitz global search in the single level case.

Still, we introduce our new simple subdivision process that is well suited to the use of DOO. The process starts with the n -dimensional hypercube (i.e. the closed ball $\mathcal{B}_{\|\cdot\|_\infty}(\mathbf{0}, 1) = \{x \in \mathbb{R}^n \mid \|x - \mathbf{0}\|_\infty \leq 1\}$ centered in $\mathbf{0}$ with a radius 1) Note that hypercubes are cells that cover $[0, 1]^n$ whereas their intersection with the unit simplex are cells covering the $n - 1$ dimensional simplex. The following theorem shows how to determine whether a hypercube intersects the simplex or not, and how to compute a *referent* point, on which the function will be evaluated. All the created hypercubes can again be covered with 2^n hypercubes, and so on, which creates an iterative covering of the unit simplex with smaller and smaller hypercubes.

THEOREM 3.4 (INTERSECTION BETWEEN THE n -DIMENSIONAL UNIT SIMPLEX AND AN n -DIMENSIONAL HYPERCUBE). *Let $n \in \mathbb{N}^* \setminus \{1\}$. Let H be an n -dimensional cube (i.e. a closed ball for $\|\cdot\|_\infty$) of radius η whose center is called m . Then, H and $S_n(1)$ intersect (i.e., $(H \cap S_n(1)) \neq \emptyset$) if and only if $\exists(x_i, x_j) \in H^2$ such that $\sum_{k=1}^n x_i^k \leq 1$ and $\sum_{k=1}^n x_j^k \geq 1$, and an intersection point (called referent point) can be computed.*

PROOF. Let us consider the diagonal from the lowest point ($x_{inf} = (m^1 - \eta, \dots, m^n - \eta)$) to the highest point ($x_{sup} = (m^1 + \eta, \dots, m^n + \eta)$) of H , and apply the Intermediate Value Theorem. There is no intersection point if $\sum_k x_{inf}^k > 1$ or $\sum_k x_{sup}^k < 1$. Else, the intersection point is $x = x_{inf} + t(x_{sup} - x_{inf})$, where $t = \frac{1 - \sum_{k=1}^n x_{inf}^k}{2n\eta}$. Indeed,

$$\sum_{k=1}^n x^k = 1 \iff \sum_{k=1}^n [x_{inf}^k + t(x_{sup}^k - x_{inf}^k)] = 1 \quad (5)$$

$$\iff t = \frac{1 - \sum_{k=1}^n x_{inf}^k}{\sum_k x_{sup}^k - \sum_k x_{inf}^k} \quad (6)$$

$$\iff t = \frac{1 - \sum_{k=1}^n x_{inf}^k}{n \cdot \|x_{sup} - x_{inf}\|_\infty} \quad (7)$$

$$\iff t = \frac{1 - \sum_{k=1}^n x_{inf}^k}{n \cdot 2\eta}. \quad (8)$$

□

Figure 3 illustrates the iterative subdivision of the 2-dimensional unit simplex by 2-dimensional hypercubes (i.e., squares). Let us point out that the subdivision operation is here concentrated around the optimum: $(0, 1)$. Squares with an “x” do not intersect with the simplex and are pruned, while the other ones do and are therefore kept. The unit-simplex being of dimensionality $n - 1$ and the hypercubes of dimensionality n , as illustrated by Figure 2, we conjecture that $\lim_{n \rightarrow \infty} \#H(N) = N \cdot 2^{n-1}$, where $\#H(N)$ denotes the number

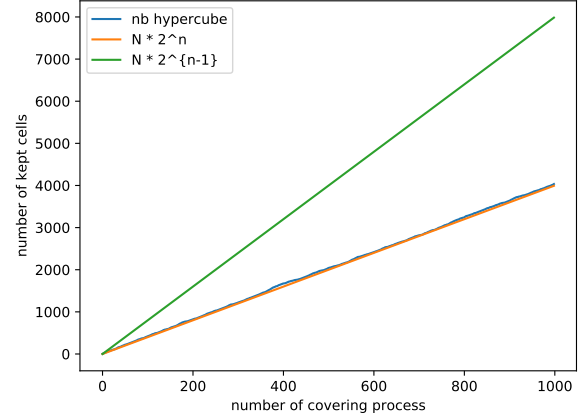


Figure 2: Number of hypercubes kept when optimizing over the n -dimensional simplex (here, $n = 3$) as a function of DOO's number of iterations (N).

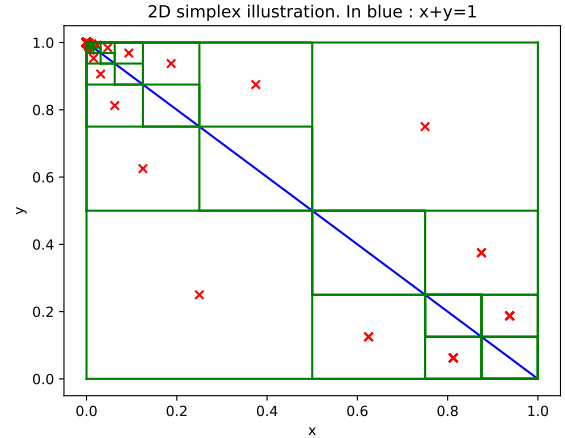


Figure 3: Illustration of the subdividing process

of hypercubes kept after N iterations and n is the dimension of the covered simplex.

4 MIN-MAX α -HÖLDER OPTIMIZATION

In this section, we come back to the min - max problem to provide a global optimization algorithm that converges in finite time to an ϵ -optimal solution.

To do so, one can first notice the α -Hölder continuity of the outer function $x \mapsto \max_{\mathbf{y}} f(x, \mathbf{y})$.

LEMMA 4.1. *Let f be a α -Hölder function and let C be one of its α -Hölder constants. The function $x \mapsto \max_{\mathbf{y}} f(x, \mathbf{y})$ is α -Hölder and C is one of its α -Hölder constants.*

PROOF. Let

$$g(\mathbf{x}) \stackrel{\text{def}}{=} \max_{\mathbf{y}} f(\mathbf{x}, \mathbf{y}).$$

Then, for any \mathbf{x} and \mathbf{x}' , assuming wlog that $g(\mathbf{x}) > g(\mathbf{x}')$,

$$\begin{aligned} g(\mathbf{x}) - g(\mathbf{x}') &= \max_{\mathbf{y}} f(\mathbf{x}, \mathbf{y}) - \max_{\mathbf{y}} f(\mathbf{x}', \mathbf{y}) \\ &\leq \max_{\mathbf{y}} (f(\mathbf{x}', \mathbf{y}) + C\|\mathbf{x}, \mathbf{y}\| - (f(\mathbf{x}', \mathbf{y}) + C\|\mathbf{x}', \mathbf{y}\|)) - \max_{\mathbf{y}} f(\mathbf{x}', \mathbf{y}) \\ &= \max_{\mathbf{y}} (f(\mathbf{x}', \mathbf{y}) + C\|\mathbf{x} - \mathbf{x}'\|^\alpha - \max_{\mathbf{y}} f(\mathbf{x}', \mathbf{y})) \\ &= C\|\mathbf{x} - \mathbf{x}'\|^\alpha. \end{aligned}$$

Similar result holds if $g(\mathbf{x}') > g(\mathbf{x})$ and this concludes the proof. \square

With this property, we solve our min - max optimization problem by using two nested DOO processes, *i.e.*,

- an outer ϵ_1 -optimal DOO minimizing the function $\mathbf{x} \mapsto \min_{\mathbf{y}} f(\mathbf{x}, \mathbf{y})$, using the solution of
- an inner ϵ_2 -optimal DOO maximizing the function $\mathbf{y} \mapsto f(\mathbf{x}, \mathbf{y})$ for fixed \mathbf{x}

(see Algorithm 2). Munos' proof can be adapted to this case to show that the final error is $\epsilon = \epsilon_1 + \epsilon_2$.

Algorithm 2: BiDOO

```

1 Fct BiDOO( $\mathcal{A}^1 \times \mathcal{A}^2 \rightarrow \mathbb{R}; x, y \mapsto f(x, y), \epsilon_1, \epsilon_2, C$ )
   input :  $f : \mathcal{A}^1 \times \mathcal{A}^2 \rightarrow \mathbb{R}$  an  $\alpha$ -Hölder function
2    $\langle x_{\min}, v_{\min} \rangle \leftarrow \text{DOO}(\left[ \mathcal{A}^1 \rightarrow \mathbb{R}; x \mapsto -\text{getMin}(\left[ \text{DOO}(\mathcal{A}^2 \rightarrow \mathbb{R}; y \mapsto -f(x, y), \epsilon_2, C) \right], \epsilon_1, C) \right])$ 
   /* getMin( $\cdot, \cdot$ ) here returns its second argument, i.e., the minimum of the inner DOO computation. */
3 return  $\langle x_{\min}, v_{\min} \rangle$ 

```

Remark 4.2 (Pruning the inner process). Using a pessimistic bound for the outer process is not only useful to provide a finite-time convergent algorithm, but it also helps to prune some irrelevant parts of the search space. The current pessimistic bound of the outer DOO is passed to the inner DOO when called, so that it can stop whenever the value of x is necessarily higher than the pessimistic bound of the inner process.

4.1 Complexity analysis

Let us now upper bound the number of iterations it takes for BiDOO to converge by analyzing it in the worst case, *i.e.*, optimizing a constant function. Let us assume that BiDOO is called to solve $(x, y) \in \mathcal{S}_n \times \mathcal{S}_p \mapsto A \in \mathbb{R}$, but is given an (overestimating) α -Hölder constant $C > 0$. In this case, BiDOO iteratively covers up the simplex with a thinner and thinner uniform grid, and the exact number of iterations it takes before reaching a stopping criterion ϵ can be found analytically.

THEOREM 4.3 (COMPLEXITY UPPER-BOUND). *For any zero-sum game $\langle \mathcal{A}^1, \mathcal{A}^2, f = (x, y) \mapsto 0 \rangle$, BiDOO return an ϵ -optimum (x^*, y^*) in less than $\left[\sum_{i=0}^{\log_2(C/\epsilon)} (2^n)^i \right] \times \left[\sum_{i=0}^{\log_2(C/\epsilon)} (2^p)^i \right]$ subdivision processes, where C is any α -Hölder constant of f given to BiDOO.*

4.2 Games with dependent feasible set

In this section, we will look at games with dependent feasible sets, formally given by the computation of

$$\min_{x \in \mathcal{A}^1} \max_{y \in F(x) \subset \mathcal{A}^2} f(x, y) \quad (9)$$

for some F . Such games received a growing interest recently, especially in the form of Fisher market problems [10], formulated as a game with dependent feasible set in [12]. Assuming that f is α -Hölder in the whole space $\mathcal{A}^1 \times \mathcal{A}^2$, the following provides a sufficient hypothesis on F to ensure BiDOO's convergence on such a game.

THEOREM 4.4. *Let $\mathcal{B}(\mathcal{A}^2)$ denote the Borelian algebra of \mathcal{A}^2 . If, $\forall x \in \mathcal{A}^1$, $F(x)$ is a bounded, closed and convex (*i.e.* compact convex) element $\mathcal{B}(\mathcal{A}^2)$ and F is λ -Lipschitz continuous for the Hausdorff distance, then BiDOO converges towards an ϵ -optimum in finite time.*

PROOF. Let us start by defining the Hausdorff distance for compact sets of $\mathcal{B}(\mathcal{A}^2)$:

$$\begin{aligned} H : \mathcal{B}(\mathcal{A}^2)^2 &\rightarrow \mathbb{R}^+ \\ (A, B) &\mapsto \max\left\{ \sup_{a \in A} \inf_{b \in B} \|a - b\|_\infty, \sup_{b \in B} \inf_{a \in A} \|a - b\|_\infty \right\}. \end{aligned}$$

Let $f : \mathcal{A}^1 \times \mathcal{A}^2 \rightarrow \mathbb{R}$ be α -Hölder, and C be a Hölder constant of f . Assume that $F : \mathcal{A}^1 \rightarrow \mathcal{B}(\mathcal{A}^2)$ is λ -Lipschitz-continuous for the Hausdorff distance.

Let $x \in \mathcal{A}^1$. Let $\epsilon > 0$. We aim at bounding the variations of $x \mapsto \max_{y \in F(x)} f(x, y)$ within a ball of radius ϵ around x , formally given by $\sup_{\tilde{x} \in B_{\|\cdot\|_\infty}(x, \epsilon)} |\max_{y \in F(x)} f(x, y) - \max_{y \in F(\tilde{x})} f(\tilde{x}, y)|$.

A first observation is that $\forall \tilde{x} \in B_{\|\cdot\|_\infty}(x, \epsilon)$,

$$\forall y \in F(\tilde{x}), f(\tilde{x}, y) \leq f(x, y) + C\|x - \tilde{x}\|_\infty^\alpha,$$

since f is α -Hölder in the whole space $\mathcal{A}^1 \times \mathcal{A}^2$, so that

$$\max_{y \in F(\tilde{x})} f(\tilde{x}, y) \leq \max_{y \in F(\tilde{x})} [f(x, y)] + C\|x - \tilde{x}\|_\infty^\alpha.$$

Now, for any $\tilde{x} \in B_{\|\cdot\|_\infty}(x, \epsilon)$, let us define I as

$$I \stackrel{\text{def}}{=} \left| \max_{y \in F(\tilde{x})} f(x, y) - \max_{y \in F(x)} f(x, y) \right|. \quad (10)$$

Let $y^* \in \arg \max_{y \in F(x)} f(x, y)$ and $\tilde{y}^* \in \arg \max_{y \in F(\tilde{x})} f(x, y)$ (whose existence is given by f 's continuity on a compact set).

Assuming wlog that $f(x, y^*) \geq f(x, \tilde{y}^*)$, we have

$$|I| = f(x, y^*) - \max_{y \in F(\tilde{x})} f(x, y) \leq f(x, y^*) - f(x, \Pi_{F(\tilde{x})}(y^*)) \quad (11)$$

where $\Pi_{F(\tilde{x})}(y) \in \arg \min_{\beta \in F(\tilde{x})} \|y - \beta\|_\infty$, which exists as a minimum as we are dealing with compact convex sets of \mathbb{R}^p . Now,

$$\begin{aligned} \|y^* - \Pi_{F(\tilde{x})}(y^*)\|_\infty &\leq \sup_{y \in F(x)} \|y - \Pi_{F(\tilde{x})}(y)\|_\infty \\ &\leq H(F(x), F(\tilde{x})). \end{aligned} \quad (13)$$

Thus,

$$f(x, y^*) - f(x, \Pi_{F(\tilde{x})}(y)) \leq CH(F(x), F(\tilde{x}))^\alpha \quad (14)$$

$$\leq C\lambda \|x - \tilde{x}\|_\infty^\alpha \quad (15)$$

and, finally, $|I| \leq C\lambda \|x - \tilde{x}\|_\infty^\alpha \leq C\lambda \epsilon^\alpha$ so that $\forall \tilde{x} \in B_{\|\cdot\|_\infty}(x, \epsilon)$,

$$\begin{aligned} \left| \min_{y \in F(x)} f(x, y) - \min_{y \in F(\tilde{x})} f(\tilde{x}, y) \right| &\leq C \|x - \tilde{x}\|_\infty^\alpha + C\lambda \|x - \tilde{x}\|_\infty \\ &\leq C\epsilon \left(1 + \frac{\epsilon^\alpha}{\epsilon} \right). \end{aligned}$$

This defines the upper bound of a cell $B_{\|\cdot\|_\infty}(x, \epsilon)$ so that DOO applies to the computation of the outer function $g : x \mapsto \max_{y \in F(x)} f(x, y)$. \square

5 EXPERIMENTS

The following aims at studying the behavior of BiDOO (i) relatively to its hyperparameters (Secs. 5.1 and 5.2) and (ii) for games with dependent feasible set (Sec. 5.3). All experiments ran on an Ubuntu machine with i7-10810U 1.10 GHz Intel processor, 16 GB available RAM and the code used is available at <https://github.com/aurelienDelageInsaLyon/minimaxDOO> under MIT license.

When considering normal-form games, the exact min - max value is given by the resolution of a linear program using Cplex². The α -Hölder constants of the inner and the outer processes are obtained using the following lemma.

LEMMA 5.1. $f : x \mapsto \max_{\mathbf{y}} \mathbf{x}^\top \cdot M \cdot \mathbf{y}$ is $\sum_j \max_i |m_{i,j}|$ -Lipschitz and $\forall x, \mathbf{y} \mapsto \mathbf{x}^\top \cdot M \cdot \mathbf{y}$ is $\sum_i \max_j |m_{i,j}|$ -Lipschitz when x, \mathbf{y} are elements of the unit simplex of \mathbb{R}^n .

PROOF. First, we show that $\forall x, \mathbf{y} \mapsto \mathbf{x}^\top \cdot M \cdot \mathbf{y}$ is $\sum_j \max_i |m_{i,j}|$ -Lipschitz.

$$\sup_{\mathbf{y}, \tilde{\mathbf{y}} \in S_n(1)} \frac{|\mathbf{x}^\top \cdot M \cdot \mathbf{y} - \mathbf{x}^\top \cdot M \cdot \tilde{\mathbf{y}}|}{\|\mathbf{y} - \tilde{\mathbf{y}}\|_\infty} \quad (16)$$

$$= \sup_{\mathbf{y}, \tilde{\mathbf{y}} \in S_n(1)} \frac{|\mathbf{x}^\top \cdot M \cdot (\mathbf{y} - \tilde{\mathbf{y}})|}{\|\mathbf{y} - \tilde{\mathbf{y}}\|_\infty} \quad (17)$$

$$= \sup_{\mathbf{y}, \tilde{\mathbf{y}} \in S_n(1)} \frac{|\sum_j (xM)_j (y_j - \tilde{y}_j)|}{\|\mathbf{y} - \tilde{\mathbf{y}}\|_\infty} \quad (18)$$

$$\leq \sup_{\mathbf{y}, \tilde{\mathbf{y}} \in S_n(1)} \frac{\sum_j |(xM)_j| |y_j - \tilde{y}_j|}{\|\mathbf{y} - \tilde{\mathbf{y}}\|_\infty} \quad (19)$$

$$\leq \sup_{\mathbf{y}, \tilde{\mathbf{y}} \in S_n(1)} \frac{\max_j |y_j - \tilde{y}_j| \sum_j |(xM)_j|}{\|\mathbf{y} - \tilde{\mathbf{y}}\|_\infty} \quad (20)$$

$$= \sup_{\mathbf{y}, \tilde{\mathbf{y}} \in S_n(1)} \sum_j \sum_i x_i |m_{i,j}| \quad (21)$$

$$\leq \sup_{\mathbf{y}, \tilde{\mathbf{y}} \in S_n(1)} \sum_j \max_i |m_{i,j}| \sum_i x_i \quad (22)$$

$$= \sum_j \max_i |m_{i,j}| \quad \text{since } \mathbf{x} \in S_n(1). \quad (23)$$

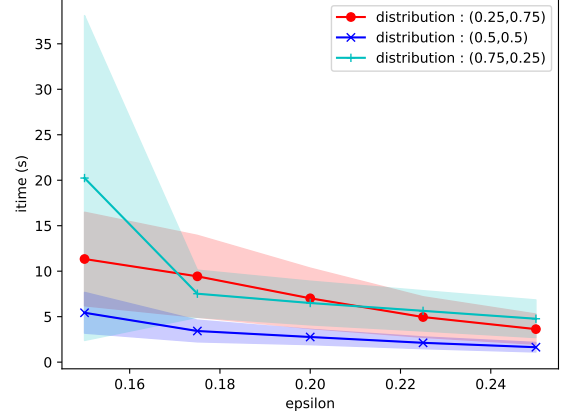


Figure 4: BiDOO's convergence time (average and standard error) as a function of ϵ for 3 different distributions.

Now, we show that $x \mapsto \max_{\mathbf{y}} \mathbf{x}^\top \cdot M \cdot \mathbf{y}$ is $\sum_i \max_j |m_{i,j}|$ -Lipschitz. $\forall(x, \tilde{x})$, assuming that $\max_{\mathbf{y}} \mathbf{x}^\top \cdot M \cdot \mathbf{y} \geq \max_{\mathbf{y}} \tilde{x}^\top \cdot M \cdot \mathbf{y}$, and noting \mathbf{y}^* an element of $\arg \max_{\mathbf{y}} \mathbf{x}^\top \cdot M \cdot \mathbf{y}$,

$$|\max_{\mathbf{y}} \mathbf{x}^\top \cdot M \cdot \mathbf{y} - \max_{\mathbf{y}} \tilde{x}^\top \cdot M \cdot \mathbf{y}| \quad (24)$$

$$\leq \mathbf{x}^\top \cdot M \cdot \mathbf{y}^* - \tilde{x}^\top \cdot M \cdot \mathbf{y}^* \quad (25)$$

$$= (\mathbf{x} - \tilde{x})^\top \cdot M \cdot \mathbf{y}^* \quad (26)$$

$$\leq \sum_i |(M \cdot \mathbf{y}^*)_i| \|x - \tilde{x}\|_\infty \quad (27)$$

$$\leq \left(\sum_i \max_j |m_{i,j}| \right) \|x - \tilde{x}\|_\infty \quad \text{since } \mathbf{y} \in S_n(1), \quad (28)$$

so that $\sup_{x, \tilde{x}} \frac{|\max_{\mathbf{y}} \mathbf{x}^\top \cdot M \cdot \mathbf{y} - \max_{\mathbf{y}} \tilde{x}^\top \cdot M \cdot \mathbf{y}|}{\|x - \tilde{x}\|_\infty} \leq \sum_i \max_j |m_{i,j}|$. \square

5.1 Choosing the ϵ -distribution

A question is whether one should choose $\epsilon_1 = \epsilon_2 = \frac{1}{2}\epsilon$ or another distribution, such as $\epsilon_1 = \frac{3}{4}\epsilon$ or $\epsilon_2 = \frac{1}{4}\epsilon$. To partially study this, Figures 4 and 5 respectively provide BiDOOs computation time and actual error when solving the min - max problem for randomly generated 3x3 matrices taken in $[0, 1]^{3 \times 3}$. Results are randomized over 30 matrixes. We observe that (i) with the 50/50 distribution, BiDOO appears to converge faster, but (ii) with the 25/75 distribution, BiDOO provides a lower error in average. While the error results appear very intuitive, the time ones are a bit more surprising. One could expect the 25/75 distribution BiDOO to be faster than the 50/50 one as it is less demanding with the inner process, which is called exponentially more often than the outer one. Still, as the number of iterations required to provably reach an error less than a given $\epsilon = \epsilon_1 + \epsilon_2$ is exponential w.r.t ϵ , the 25/75 distribution BiDOO may require way more time for the outer DOO to converge. This might, as a whole, take more time. In all the following experiments, we pick $\epsilon_1 = \epsilon_2$.

²Cplex 12.1 free version <https://www.ibm.com/fr-fr/analytcs/cplex-optimizer>

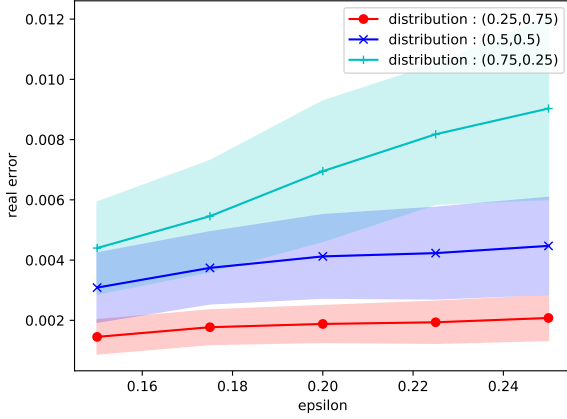


Figure 5: BiDOO’s observed error (average and standard error) as a function of ϵ for 3 different distributions.

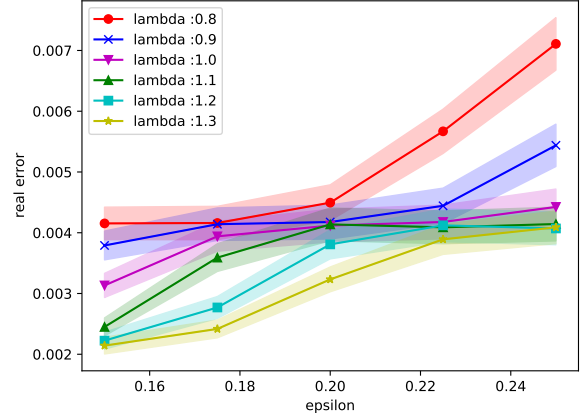


Figure 7: Error of Algorithm 2 compared to the exact min - max value given by Cplex as a function of the error ϵ imposed, for different Hölder constants. If C_{ext} and C_{int} respectively denote the 1-Hölderian constants of the outer and the inner DOO processes, the curve labelled “lambda : 1.1” corresponds to launching BiDOO with $C_{ext} \cdot 1.1$ and $C_{int} \cdot 1.1$.

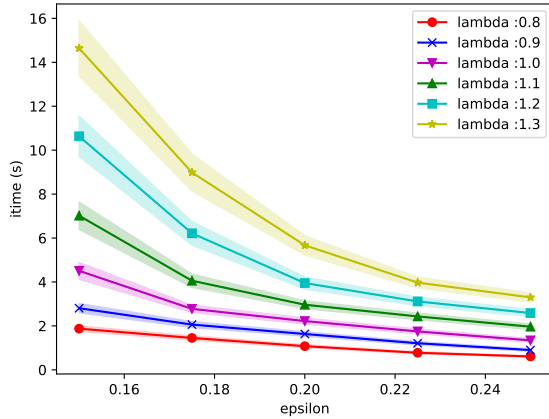


Figure 6: Computation time (in seconds) for Algorithm 2 as a function of the error ϵ imposed, for different Hölder constants. If C_{ext} and C_{int} respectively denote the 1-Hölderian constants of the outer and the inner DOO processes, the curve labeled “lambda : 1.1” corresponds to launching BiDOO with $C_{ext} \cdot 1.1$ and $C_{int} \cdot 1.1$.

5.2 Validating the approach

We consider 100 randomly generated matrices $M \in [0, 1]^{3 \times 3}$. Figures 6 and 7 respectively give the error compared to the exact value and the computation time as a function of the error imposed to the algorithm.

Several things are to be noted on this figure. Firstly, we observe consistency, *i.e.* BiDOO’s error with respect to the exact value always being lower than the error ϵ imposed to DOO. Secondly, and as predicted by Section 4.1, the time taken by the algorithm to converge appears exponential with respect to ϵ . Interestingly, DOO returns a value that is guaranteed to be ϵ -close to the optimal, but

which is actually around $\epsilon/10$ -close to the exact value returned by linear programming.

Looking back at Theorem 4.3, we compared a uniform search iteratively building a small enough uniform grid to BiDOO using the same set of 100 randomly-generated 3×3 matrices, for $\epsilon = 0.15$. It takes the uniform search in average 7.8 s (std. error 0.9 s) to end, whereas BiDOO converges on average in 2.4 s (std. error 0.2 s).

5.3 Comparison with the state of the art

Bilinear Games and Gradient Descent Ascent. It is known [30] that “vanilla” GDA (*i.e.* with constant step size) can fail even for the bilinear game $x^T \cdot I_n \cdot y$, where $x, y \in \mathbb{R}^n$, and I_n is the n -dimensional identity matrix. An illustration of GDA’s behavior is given in Figure 8. As the game is bilinear, it is *a fortiori* α -Hölder and BiDOO provably solves the min - max problem.

Games with dependent feasible sets. Goktas and Greenwald [13] illustrate the difficulties with running a gradient-based algorithm to solve general min - max problems through the computation of

$$\min_{x \in [-1, 1]} \max_{y \in [-1, 1], 1 - (x+y) \geq 0} x^2 + y + 1. \quad (29)$$

In such a game, GDA converges to $(0, 1)$, which is not a solution of the min - max problem, whereas BiDOO does converge as illustrated in Figure 9, provided the subdivision process is modified to satisfy the constraint $y \in [-1, 1], 1 - (x+y) \geq 0$. This result is expected as this example satisfies the assumption of Section 4.2 as proven below.

Example 5.2. $F : x \mapsto \{y | 1 - (x+y) \geq 0, y \geq 0\}$ only yields compact convex subsets of \mathbb{R} , and is Lipschitz for the Hausdorff distance H .

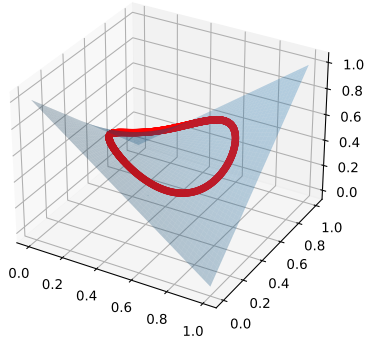


Figure 8: GDA applied to the problem $\max_x \min_y x \cdot y^\top$. Red points correspond to points visited by the GDA algorithm.

PROOF. Let $x \in [0, 1]$. We have $F(x) = \{y | 1 - (x + y) \geq 0, y \geq 0\} = \{y | y \leq 1 - x, y \geq 0\} = [0, 1 - x]$, which is bounded, closed (*i.e.* compact) and convex. Now, for any $(x, \tilde{x}) \in [0, 1]^2$,

$$H(F(x), F(\tilde{x})) = \max \left\{ \sup_{y \in [0, 1-x]} \inf_{\tilde{y} \in [0, 1-\tilde{x}]} |y - \tilde{y}|, \right. \quad (30)$$

$$\left. \sup_{\tilde{y} \in [0, 1-\tilde{x}]} \inf_{y \in [0, 1-x]} |y - \tilde{y}| \right\}. \quad (31)$$

Assume now without loss of generality that $x > \tilde{x}$. Then, $F(x) = [0, 1 - x] \subset [0, 1 - \tilde{x}] = F(\tilde{x})$ and

$$\forall \tilde{y} \in F(\tilde{x}), \inf_{y \in F(x)} |y - \tilde{y}| = \begin{cases} 0 & \text{if } \tilde{y} \in F(x), \\ |\tilde{y} - (1 - x)| & \text{else.} \end{cases}$$

Thus,

$$\sup_{\tilde{y} \in F(\tilde{x})} \inf_{y \in F(x)} |y - \tilde{y}| = |1 - \tilde{x} - (1 - x)|.$$

Now, since $F(x) \subset F(\tilde{x})$,

$$\forall y \in F(x), \inf_{\tilde{y} \in F(\tilde{x})} |y - \tilde{y}| = |y - y| = 0, \text{ so that}$$

$$\sup_{y \in F(x)} \inf_{\tilde{y} \in F(\tilde{x})} |y - \tilde{y}| = 0.$$

Finally, $H(F(x), F(\tilde{x})) = |x - \tilde{x}|$ and F is 1-Lipschitz, thus justifying that BiDOO successfully solves the min - max problem of Section 4.2. \square

Interestingly, Figure 9 shows a pessimistic bound within ϵ to the optimum way earlier than the optimistic bound. This behavior is frequent with approaches which guarantee ϵ -optimality using optimistic and pessimistic bounds; proving optimality (*i.e.*, lowering the optimistic bound ϵ -closely to the optimum) is harder than heuristically finding near-optimal points.

6 DISCUSSION

The time complexity of the proposed algorithm was shown to be exponential with respect to (i) the dimension of the game and (ii) $1/\epsilon$. This theoretical claim is supported by our experiments. This is an expected result as computing local Nash Equilibrium,

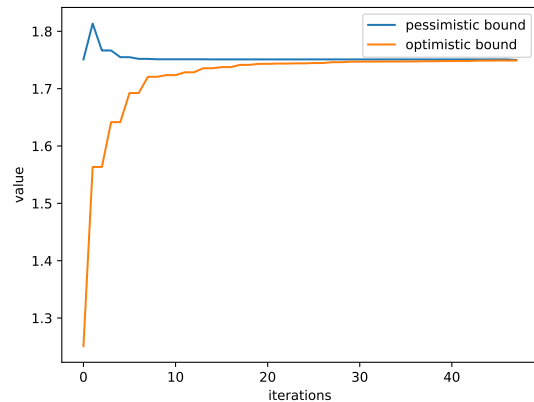


Figure 9: Upper and lower bounds as a function of the iteration number.

when the game has stronger properties than the α -Hölder one (differentiability and Lipschitz gradient) [6], is already in ExpTime with respect to either the game's dimension, $1/\epsilon$ or the game's smoothness. We introduced the BiDOO algorithm, which is (to the best of our knowledge) the first algorithm to provably converge towards an ϵ -min - max of any α -Hölder game. The algorithm being simple, we believe that it may inspire further work. As an example, if the game's payoff is convex with respect to player 2, then the inner DOO process can be replaced by a quicker gradient descent. Moreover, our algorithm can be extended to the unknown smoothness case, building on SOO [19] instead of DOO. BiDOO might also pave the way to tackle general-sum games with any number of players with independent divergent payoff functions by extending our algorithm to multi-level optimization problems. Still, this remains an open question as the function giving 1's value for a strategy x given that 2 best responds to it according to her payoff function is discontinuous, and the computation of discontinuities is a harder problem than computing a Stackelberg Equilibrium. Showing that there only exists a finite number of discontinuities (which was proven for the one-dimensional case for a dense set of payoff functions in by Bressan and Jiang [3]) should enable the use of an ϵ -greedy version of BiDOO asymptotically converging towards an ϵ -Stackelberg equilibrium for any α -Hölder game.

7 CONCLUSION

This paper proposes a method to deal with min - max optimization of α -Hölder functions, leading to a finite-time convergent algorithm towards an ϵ -global optimum. We focused on the optimization over simplices, derived an algorithm and evaluated it. We hope our method can be useful whenever such a problem arises.

REFERENCES

- [1] Eugene Asarin, Julien Cervelle, Aldric Degorre, Catalin Dima, Florian Horn, and Victor Kozyakin. Entropy games and matrix multiplication games. *arXiv preprint arXiv:1506.04885*, 2015.
- [2] Sarah Brandsen, Isabelle Jianing Geng, and Gilad Gour. What is entropy? A perspective from games of chance. *Physical Review E*, 105(2):024117, 2022.

- [3] Alberto Bressan and Yilun Jiang. On the generic structure and stability of stackelberg equilibria. *Journal of Optimization Theory and Applications*, 183: 840–880, 2019.
- [4] Nicolas Brunner and Noah Linden. Connection between Bell nonlocality and Bayesian game theory. *Nature communications*, 4(1):1–6, 2013.
- [5] Yinlam Chow, Aviv Tamar, Shie Mannor, and Marco Pavone. Risk-sensitive and robust decision-making: A CVaR optimization approach. *Advances in neural information processing systems*, 28, 2015.
- [6] Constantinos Daskalakis. Non-concave games: A challenge for game theory's next 100 years. *Nobel symposium "One Hundred Years of Game Theory: Future Applications and Challenges*, 2021.
- [7] Constantinos Daskalakis, Stratis Skoulakis, and Manolis Zampetakis. The complexity of constrained min-max optimization. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 1466–1478, 2021.
- [8] Jianbo Du, Liqiang Zhao, Jie Feng, and Xiaoli Chu. Computation offloading and resource allocation in mixed fog/cloud computing systems with min-max fairness guarantee. *IEEE Transactions on Communications*, 66(4):1594–1608, 2017.
- [9] Herbert Edelsbrunner and Daniel R Grayson. Edgewise subdivision of a simplex. In *Proceedings of the fifteenth annual symposium on Computational geometry*, pages 24–30, 1999.
- [10] Irving Fisher. *Mathematical investigations in the theory of value and prices, and appreciation and interest*. PhD thesis, Yale University, 1892.
- [11] Daniel Friedman. On economic applications of evolutionary game theory. *Journal of Evolutionary Economics*, 8(1):15–43, 1998.
- [12] Denizalp Goktas and Amy Greenwald. Convex-concave min-max Stackelberg games. *Advances in Neural Information Processing Systems*, 34:2991–3003, 2021.
- [13] Denizalp Goktas and Amy Greenwald. Gradient descent ascent in min-max Stackelberg games. *arXiv preprint arXiv:2208.09690*, 2022.
- [14] Matthew Jagielski, Michael Kearns, Jieming Mao, Alina Oprea, Aaron Roth, Saeed Sharifi-Malvajerdi, and Jonathan Ullman. Differentially private fair learning. In *International Conference on Machine Learning*, pages 3000–3008. PMLR, 2019.
- [15] Chi Jin, Praneeth Netrapalli, and Michael I. Jordan. What is local optimality in nonconvex-nonconcave minimax optimization?, 2019.
- [16] Tianyi Lin, Chi Jin, and Michael Jordan. On gradient descent ascent for nonconvex-concave minimax problems. In *International Conference on Machine Learning*, pages 6083–6093. PMLR, 2020.
- [17] Tung Mai, Milena Mihail, Ioannis Panageas, Will Ratcliff, Vijay Vazirani, and Peter Yunker. Cycles in zero-sum differential games and biological diversity. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pages 339–350, 2018.
- [18] Lars Mescheder, Andreas Geiger, and Sebastian Nowozin. Which training methods for GANs do actually converge? 2018. doi: 10.48550/ARXIV.1801.04406.
- [19] Rémi Munos. Optimistic optimization of a deterministic function without the knowledge of its smoothness. *Advances in Neural Information Processing Systems*, 24, 2011.
- [20] Rémi Munos. From bandits to Monte-Carlo Tree Search: The optimistic principle applied to optimization and planning. *Foundations and Trends in Machine Learning*, 7(1):1–130, 2014.
- [21] Maher Nouiehed, Maziar Sanjabi, Tianjian Huang, Jason D Lee, and Meisam Razaviyayn. Solving a class of non-convex min-max games using iterative first order methods. *Advances in Neural Information Processing Systems*, 32, 2019.
- [22] Frans A Oliehoek, Rahul Savani, Jose Gallego-Posada, Elise Van der Pol, Edwin D De Jong, and Roderich Groß. GANGs: Generative adversarial network games. *arXiv preprint arXiv:1712.00679*, 2017.
- [23] Remigijus Paulavičius and Julius Žilinskas. *Simplicial global optimization*. Springer, 2014.
- [24] Hassan Rafique, Mingrui Liu, Qihang Lin, and Tianbao Yang. Weakly-convex concave min-max optimization: Provable algorithms and applications in machine learning. 2018. doi: 10.48550/arxiv.1810.02060.
- [25] Maziar Sanjabi, Jimmy Ba, Meisam Razaviyayn, and Jason D Lee. On the convergence and robustness of training GANs with regularized optimal transport. *Advances in Neural Information Processing Systems*, 31, 2018.
- [26] Prasanna Sattigeri, Samuel C Hoffman, Vijil Chenthamarakshan, and Kush R Varshney. Fairness GAN. *arXiv preprint arXiv:1805.09910*, 2018.
- [27] Yoav Shoham and Kevin Leyton-Brown. *Multiagent systems: Algorithmic, game-theoretic, and logical foundations*. Cambridge University Press, 2008.
- [28] Maurice Sion. On general minimax theorems. *Pacific Journal of mathematics*, 8(1):171–176, 1958.
- [29] John von Neumann. Zur Theorie der Gesellschaftsspiele. *Mathematische Annalen*, 100, 1928. URL <https://doi.org/10.1007/BF01448847>.
- [30] Jiawei Zhang, Peijun Xiao, Ruoyu Sun, and Zhiquan Luo. A single-loop smoothed gradient descent-ascent algorithm for nonconvex-concave min-max problems. *Advances in Neural Information Processing Systems*, 33:7377–7389, 2020.