Czech Technical University in Prague
Faculty of Electrical Engineering
Department of Cybernetics

Computing Minimax Strategies in Nonconvex-Nonconcave Strategic Games

Bachelor’s thesis of
Marika Kosohorská

Supervised by
doc. Ing. Tomáš Kroupa, Ph.D.

May 2024

ABSTRACT

This thesis explores various concepts extending Nash equilibrium in two-player zero-sum games with compact strategy spaces and possibly nonconvex-nonconcave loss functions. We discuss the theoretical characteristics of local and approximate Nash equilibria, as well as their optimality conditions. The main focus of the thesis is on the min-max critical point, which is a first-order solution concept extending Nash equilibrium. We prove some characteristics of the min-max critical point. We propose the implementation of the StayOnTheRidge algorithm for finding min-max critical points in Julia and compare the results of the algorithm with the results obtained by other algorithms on various examples. We also present an extension of the STON'R to hyperrectangle and discuss a general challenge of the algorithm's modification to operate on the cartesian product of simplices. The theoretical result of this thesis is the introduction of the concept of the generalized min-max critical point, which extends the min-max critical point to locally Lipschitz functions. We prove the existence of the solution to the corresponding generalized variational inequality and show some properties of the generalized min-max critical points.

KEYWORDS: Two-player zero-sum games, Nash equilibrium, Min-max critical point, Variational inequality, Clarke analysis

ABSTRAKT

Tato práce zkoumá různé koncepty Nashových ekvilibrií dvouhráčových her s nulovým součtem a kompaktními prostory strategií, přičemž ztrátové funkce mohou být nekonvexně-nekonkávní. Diskutujeme teoretické vlastnosti lokálních a aproximovaných Nashových ekvilibrií a jejich podmínky optimality. Hlavním zaměřením práce je koncept min-max kritického bodu, který představuje řešení prvního řádu rozšiřující Nashovo ekvilibrium. Dokazujeme některé charakteristiky min-max kritického bodu. Představujeme implementaci algoritmu StayOnTheRidge pro nalezení min-max kritického bodu v jazyce Julia a porovnáváme kvalitu výsledků s jinými algoritmy na různých příkladech. Dále popisujeme rozšíření algoritmu STON'R na hyperobdélník a je nastíněna výzva zobecnění algoritmu na obecnější množiny, jako je kartézský součin simplexů. Teoretickým výsledkem této práce je zavedení pojmu zobecněného min-max kritického bodu, který rozšiřuje koncept min-max kritického bodu na lokálně Lipschitzovské funkce. Dokazujeme existenci řešení odpovídající zobecněné variační nerovnice a ukazujeme některé vlastnosti zobecněných min-max kritických bodů.

KEYWORDS: Dvouhráčové hry s nulovým součtem, Nashovo ekvilibrium, Min-max kritický bod, Variační nerovnice, Clarkova analýza

BACHELOR‘S THESIS ASSIGNMENT

I. Personal and study details

Student's name:	**Kosohorská Marika**	Personal ID number:	**507284**
Faculty / Institute:	**Faculty of Electrical Engineering**		
Department / Institute:	**Department of Cybernetics**		
Study program:	**Open Informatics**		
Specialisation:	**Artificial Intelligence and Computer Science**		

II. Bachelor’s thesis details

Bachelor’s thesis title in English:

Computing Minimax Strategies in Nonconvex-Nonconcave Strategic Games

Bachelor’s thesis title in Czech:

Výpočet minimaxových strategií v nekonvexně-nekonkávních hrách

Guidelines:

Nonconvex-nonconcave games represent a challenging family of games [2]. Unlike convex-concave games where the objectives for all players are either convex or concave, nonconvex-nonconcave games lead to more unpredictable equilibria. This complexity often results in multiple local minima and maxima. Additionally, the lack of convexity and concavity undermines the use of traditional optimization methods that rely on gradient-based techniques. This necessitates the development of more sophisticated algorithms capable of navigating the intricate topology of nonconvex-nonconcave games [1]. As a result, solving these games requires advanced mathematical tools (non-monotone variational inequalities, gradient descent-ascent techniques for nonconvex problems) and a deeper understanding of dynamic systems, making them a formidable challenge in both theoretical and practical applications such as GAN training [4] or adversarial ML problems in general [3]. The objectives of this bachelor thesis are outlined as follows.

1. To develop further the proof of concept implementation of the STON’R (Stay-on-the-Ridge) algorithm (https://gitlab.fel.cvut.cz/kosohmar/StayOnTheRidge.jl) for computing min-max critical points.
2. To test and debug the implementation of STON’R using existing examples of continuous strategic games in literature.
3. To compare the performance of this implementation to existing iterative methods for finding Nash equilibria (Double Oracle algorithms, GDA etc.)
4. (Optional and challenging.) To extend the applicability of STON’R to possibly non-smooth utility functions or to more complex strategy sets than hypercubes.

Bibliography / sources:

[1] Daskalakis, Constantinos, et al. "Stay-on-the-ridge: Guaranteed convergence to local minimax equilibrium in nonconvex-nonconcave games." The Thirty Sixth Annual Conference on Learning Theory. PMLR, 2023.
[2] Daskalakis, Constantinos. "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.
[3] Razaviyayn, Meisam, et al. "Nonconvex min-max optimization: Applications, challenges, and recent theoretical advances." IEEE Signal Processing Magazine 37.5 (2020): 55-66.
[4] Farnia, Farzan, and Asuman Ozdaglar. "Do GANs always have Nash equilibria?." International Conference on Machine Learning. PMLR, 2020.

Name and workplace of bachelor's thesis supervisor:

doc. Ing. Tomáš Kroupa, Ph.D. Artificial Intelligence Center FEE

Name and workplace of second bachelor's thesis supervisor or consultant:

Date of bachelor's thesis assignment: **21.01.2024** Deadline for bachelor thesis submission: **24.05.2024**

Assignment valid until: **21.09.2025**

doc. Ing. Tomáš Kroupa, Ph.D. Supervisor's signature	prof. Dr. Ing. Jan Kybic Head of department's signature	prof. Mgr. Petr Páta, Ph.D. Dean's signature

III. Assignment receipt

The student acknowledges that the bachelor's thesis is an individual work. The student must produce her thesis without the assistance of others, with the exception of provided consultations. Within the bachelor's thesis, the author must state the names of consultants and include a list of references.

.

Date of assignment receipt	Student's signature

DECLARATION

I declare that the presented work was developed independently and that I have listed all sources of information used within it in accordance with the methodical instructions for observing the ethical principles in the preparation of university theses.

Prague, May 24, 2024

ACKNOWLEDGMENTS

I would like extend my sincere thanks to doc. Ing. Tomáš Kroupa, Ph.D., for introducing me to interesting problems in game theory and for patiently supervising the thesis.

CONTENTS

1 INTRODUCTION

Identifying the Nash equilibrium is a significant challenge in game theory. It can be difficult even for two-player zero-sum games with compact strategy spaces when the loss function is not convex-concave. Although the mixed Nash equilibrium is guaranteed to exist under the continuity of the loss function, the pure Nash equilibrium may not exist even for simply behaved functions. And if it exists, finding it is NP-hard [1].

Therefore, first-order solution concepts of Nash equilibrium have been introduced, referred to as min-max critical point [2], first-order Nash equilibrium [3], [4], or game-stationary solution [5]. These points represent solutions to a specific non-monotone variational inequality and are assured to exist for smooth functions. They correspond to the fixed points of the projected gradient descent-ascent dynamics and hence their computational complexity lies in the PPAD [6] class.

The STON'R algorithm [2] for finding min-max critical points is based on a topological argument, leveraging the equivalence between the fixed points of the projected gradient descent-ascent dynamics and min-max critical points. Other algorithms for finding first-order solutions have also been proposed [3], [4]. This is because finding (relaxed) Nash equilibria has practical applications in various fields such as signal and data processing, GANs training, and robust machine learning, to mention only a few [5].

1.1 SUMMARY OF THE THESIS

- This thesis examines various notions extending Nash equilibrium, such as local Nash equilibrium, approximate local Nash equilibrium, or the first order-solution concept named min-max critical point. We formulate and show some properties of the min-max critical points (Proposition 7, Proposition 9).

- Since the loss functions are not smooth in some applications, we extend the notion of the min-max critical point and the first-order Nash equilibrium to locally Lipschitz functions (Definition 14, Definition 15) and prove that the corresponding generalized variational inequality has a solution (Proposition 10). We also show some properties of the generalized min-max critical points (Proposition 12).

- We develop the implementation of the STON'R algorithm in Julia and demonstrate its results on various examples.

We also formulate an extension of the algorithm to the general hyperrectangle and outline the challenge of the algorithm's extension to the cartesian product of simplices.

- The numerical experiments compare the quality of the results produced by three algorithms: STON'R [2], RNI-SGD [3], and DO [7]. We measure the quality of the solutions by exploitability, which is defined as the sum of the differences between each player's strategy and their optimal strategy.

2 STRATEGIC GAMES

In this chapter, we focus on solution concepts for two-player zero-sum games. These games involve two players with directly opposing interests, so the total sum of their losses is always zero. One common example is the rock-paper-scissors game, represented by the following matrix.

	Rock	Paper	Scissors
Rock	0	1	-1
Paper	-1	0	1
Scissors	1	-1	0

Table 2.1: Rock-paper-scissors game matrix

The table shows the losses for the row player. Each player chooses from three pure strategies. The optimal mixed strategy for both players is to select each option with a probability $1/3$.

Nevertheless, our focus will shift towards more complex games where players have an infinite number of pure strategies to choose from. We call these games *continuous games* and formally define them in the following section.

2.1 CONTINUOUS GAMES

Min-player and Max-player select strategies $\boldsymbol{\theta} = (x_1, \dots, x_{d_1}) \in \Theta$ and $\boldsymbol{\omega} = (x_{d_1+1}, \dots, x_{d_1+d_2}) \in \Omega$, where $\Theta \subseteq \mathbb{R}^{d_1}$ and $\Omega \subseteq \mathbb{R}^{d_2}$ are compact convex sets. Let $n = d_1 + d_2$ be the dimension of the game and $K = \Theta \times \Omega \subseteq R^n$ be the joint strategy space. The loss function for Min-player, denoted by $f : K \to \mathbb{R}$, is continuous (unless otherwise stated). The loss function for Max-player is $-f$. We call such strategic game a *continuous game*. [1] The function f may not be convex-concave, i.e., it may fail to be convex in $\boldsymbol{\theta}$ for some $\boldsymbol{\omega}$, or it may fail to be concave in $\boldsymbol{\omega}$ for some $\boldsymbol{\theta}$.

For example, viewing rock-paper-scissors as a continuous game, players select their strategies $\boldsymbol{\theta} = (x_1, x_2, x_3)$ and $\boldsymbol{\omega} = (x_4, x_5, x_6)$ from 2-dimensional simplices with vertices corresponding to the rock,

1 To simplify, we define the game as continuous. We will introduce additional assumptions for f as needed throughout this thesis, although continuity will be assumed in all cases.

paper, and scissors strategies, respectively. The loss function for Min-player is a multilinear function

$$f(\boldsymbol{\theta}, \boldsymbol{\omega}) = \boldsymbol{\theta}^{\top} \begin{pmatrix} 0 & 1 & -1 \\ -1 & 0 & 1 \\ 1 & -1 & 0 \end{pmatrix} \boldsymbol{\omega}.$$

The pure strategy Nash equilibrium exists for this game and corresponds to the $\boldsymbol{\theta}^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3}) = \boldsymbol{\omega}^*$.

2.2 NASH EQUILIBRIA

In this section, we describe various notions of Nash equilibria for continuous games and discuss their properties. Nash equilibrium is a stable state in which no player has motivation to unilaterally change their strategy.

2.2.1 *Global Nash equilibria*

The global Nash equilibrium is a solution concept where both players adopt strategies that are optimal over their entire strategy spaces. Let us start with the definition.

Also called saddle point.

Definition 1 The *Nash equilibrium (NE)* of a continuous game is a pair $(\boldsymbol{\theta}^*, \boldsymbol{\omega}^*) \in \Theta \times \Omega$ such that

$$f(\boldsymbol{\theta}^*, \boldsymbol{\omega}) \leqslant f(\boldsymbol{\theta}^*, \boldsymbol{\omega}^*) \leqslant f(\boldsymbol{\theta}, \boldsymbol{\omega}^*)$$

holds for all $\boldsymbol{\theta} \in \Theta$ and $\boldsymbol{\omega} \in \Omega$.

This definition implies that $\boldsymbol{\theta}^*$ is a global minimum of $f(\cdot, \boldsymbol{\omega}^*)$ for fixed $\boldsymbol{\omega}^*$, and $\boldsymbol{\omega}^*$ is a global maximum of $f(\boldsymbol{\theta}^*, \cdot)$ for fixed $\boldsymbol{\theta}^*$.

We outline two characterizations of Nash equilibria.

Proposition 1 The continuous game has a Nash equilibrium if, and only if,

$$\min_{\boldsymbol{\theta} \in \Theta} \max_{\boldsymbol{\omega} \in \Omega} f(\boldsymbol{\theta}, \boldsymbol{\omega}) = \max_{\boldsymbol{\omega} \in \Omega} \min_{\boldsymbol{\theta} \in \Theta} f(\boldsymbol{\theta}, \boldsymbol{\omega}).$$

If f has a NE, it corresponds to the solution to $\min\limits_{\boldsymbol{\theta} \in \Theta} \max\limits_{\boldsymbol{\omega} \in \Omega} f(\boldsymbol{\theta}, \boldsymbol{\omega})$.

Proposition 1 implies that the set of solutions to $\min\limits_{\boldsymbol{\theta} \in \Theta} \max\limits_{\boldsymbol{\omega} \in \Omega} f(\boldsymbol{\theta}, \boldsymbol{\omega})$ and the set of solutions to $\max\limits_{\boldsymbol{\omega} \in \Omega} \min\limits_{\boldsymbol{\theta} \in \Theta} f(\boldsymbol{\theta}, \boldsymbol{\omega})$ are either identical or have an empty intersection. The second characterization needs the notion of exploitability.

Definition 2 Define *exploitability* $e(\mathbf{x}^*) \in \mathbb{R}$ at point $\mathbf{x}^* = (\boldsymbol{\theta}^*, \boldsymbol{\omega}^*) \in K$ as the sum of differences between each player's strategy and their optimal strategy:

$$
\begin{aligned}
e(\mathbf{x}^*) &= \max_{\boldsymbol{\theta} \in \Theta} (-f(\boldsymbol{\theta}, \boldsymbol{\omega}^*)) - (-f(\boldsymbol{\theta}^*, \boldsymbol{\omega}^*)) + \max_{\boldsymbol{\omega} \in \Omega} f(\boldsymbol{\theta}^*, \boldsymbol{\omega}) - f(\boldsymbol{\theta}^*, \boldsymbol{\omega}^*) \\
&= \max_{\boldsymbol{\omega} \in \Omega} f(\boldsymbol{\theta}^*, \boldsymbol{\omega}) - \min_{\boldsymbol{\theta} \in \Theta} f(\boldsymbol{\theta}, \boldsymbol{\omega}^*).
\end{aligned}
$$

Proposition 2 Let $\mathbf{x}^* = (\boldsymbol{\theta}^*, \boldsymbol{\omega}^*) \in K$. A point $\mathbf{x}^*$ is a Nash equilibrium if, and only if, $e(\mathbf{x}^*) = 0$.

Proof $\Rightarrow$ The right implication is trivial. From the definition of NE, we have $\max_{\boldsymbol{\theta} \in \Theta}(-f(\boldsymbol{\theta}, \boldsymbol{\omega}^*)) + \max_{\boldsymbol{\omega} \in \Omega} f(\boldsymbol{\theta}^*, \boldsymbol{\omega}) = -f(\mathbf{x}^*) + f(\mathbf{x}^*) = 0$.
$\Leftarrow$ We know that $e(\mathbf{x}^*) \geqslant 0$. Assuming the point $\mathbf{x}^*$ is not a NE, at least one player does not play optimally. This implies that either $\max_{\boldsymbol{\theta} \in \Theta}(-f(\boldsymbol{\theta}, \boldsymbol{\omega}^*)) - (-f(\mathbf{x}^*)) > 0$, or $\max_{\boldsymbol{\omega} \in \Omega} f(\boldsymbol{\theta}^*, \boldsymbol{\omega}) - f(\mathbf{x}^*) > 0$, or both. Hence, $e(\mathbf{x}^*) > 0$.

The concept of exploitability will be essential in Chapter 4, where we will compare it across solutions obtained from different methods.

If f is convex-concave, the existence of Nash equilibrium is guaranteed [8]. Recall that our area of study is the scenario when f may fail to be convex-concave. This setting yields substantial challenges, as the existence of Nash equilibria is not assured, and determining whether such a point exists is an NP-hard problem [1]. Nash equilibrium may not exist even for well-behaved functions; one such example is a convex function $f(\theta, \omega) := (\theta - \omega)^2$ over $\Theta = \Omega = [0, 1]$.

The generalization of Nash equilibrium is a mixed strategy Nash equilibrium. The concept of mixed strategy allows every player to randomize with respect to any probability measure on their strategy set. We will define concepts exclusively for Min-player; the definitions for Max-player are analogous.

Definition 3 The set of all *mixed strategies* for Min-player is defined as the set of Borel probability measures over Θ, denoted by $\Delta(\Theta)$.

The mixed strategies allows us to associate each pure strategy $\boldsymbol{\theta} \in \Theta$ with a Dirac measure from $\Delta(\Theta)$. Put $\Delta := \Delta(\Theta) \times \Delta(\Omega)$. If players adopt a *mixed strategy profile* $(p, q) \in \Delta$, the expected loss $L \colon \Delta \to \mathbb{R}$ for Min-player is defined as

$$
L(p, q) := \int_{\Theta \times \Omega} f \, d(p \times q).
$$

The mixed strategy Nash equilibrium occurs when no player wants to change their mixed strategy. In continuous games, its existence is guaranteed by Glicksberg's theorem [9].

Definition 4 A mixed strategy profile $(p^*, q^*) \in \Delta$ is a *mixed strategy Nash equilibrium* if

$$L(p^*, q) \leqslant L(p^*, q^*) \leqslant L(p, q^*) \tag{1}$$

holds for all $(p, q) \in \Delta$.

2.2.2 *Local Nash equilibria*

As mentioned, finding Nash equilibrium is a very difficult problem. This has led to the study of the local/approximate solutions. One such notion is the local Nash equilibrium [10].

Definition 5 Let $\delta > 0$. The point $(\boldsymbol{\theta}^*, \boldsymbol{\omega}^*) \in \Theta \times \Omega$ is called a *local Nash equilibrium* if it satisfies:

$$\begin{aligned} f(\boldsymbol{\theta}^*, \boldsymbol{\omega}^*) < f(\boldsymbol{\theta}, \boldsymbol{\omega}^*) &\quad \text{for all } \boldsymbol{\theta} \in \Theta \text{ such that } \|\boldsymbol{\theta} - \boldsymbol{\theta}^*\| \leqslant \delta, \\ f(\boldsymbol{\theta}^*, \boldsymbol{\omega}^*) > f(\boldsymbol{\theta}^*, \boldsymbol{\omega}) &\quad \text{for all } \boldsymbol{\omega} \in \Omega \text{ such that } \|\boldsymbol{\omega} - \boldsymbol{\omega}^*\| \leqslant \delta. \end{aligned}$$

In this definition, each player is required to play optimally only within the δ neighborhood around their selected strategy.

We formulate the optimality conditions for a local NE. Let us focus on the interior of K first, assuming that it is nonempty [10].

Proposition 3 (First-order necessary condition) If f is differentiable, any local Nash equilibrium $\mathbf{x}^* \in \text{int}(K)$ satisfies $\nabla_\Theta f(\mathbf{x}^*) = \mathbf{0}$ and $\nabla_\Omega f(\mathbf{x}^*) = \mathbf{0}$.

Proposition 4 (Second-order necessary condition) Assuming that f is twice-differentiable, any local Nash equilibrium $\mathbf{x}^* \in \text{int}(K)$ satisfies $\nabla^2_{\Theta\Theta} f(\mathbf{x}^*) \succeq 0$ and $\nabla^2_{\Omega\Omega} f(\mathbf{x}^*) \preceq 0$.

Proposition 5 (Second-order sufficient condition) Assuming that f is twice-differentiable, any stationary point $\mathbf{x}^* \in \text{int}(K)$ satisfying the following condition is a local Nash equilibrium: $\nabla^2_{\Theta\Theta} f(\mathbf{x}^*) \succ 0$ and $\nabla^2_{\Omega\Omega} f(\mathbf{x}^*) \prec 0$.

We now outline the first-order optimality condition for the entire set K [11].

Definition 6 We call a set $N_C(\mathbf{x}) = \{\mathbf{d} \in \mathbb{R}^m \mid \langle \mathbf{d}, \mathbf{x} - \mathbf{x}' \rangle \geqslant 0, \forall \mathbf{x}' \in C\}$ the *normal cone* at point $\mathbf{x} \in C$, where C is a compact convex subset of $\mathbb{R}^m$.

Proposition 6 (First-order necessary condition) Assuming f is differentiable, any local Nash equilibrium $\mathbf{x}^* = (\boldsymbol{\theta}^*, \boldsymbol{\omega}^*) \in K$ satisifes $-\nabla_\Theta f(\mathbf{x}^*) \in N_\Theta(\boldsymbol{\theta}^*)$ and $\nabla_\Omega f(\mathbf{x}^*) \in N_\Omega(\boldsymbol{\omega}^*)$.

In the interior of K, the condition reduces to the condition in Proposition 3 because the normal cone at any point within int(K) contains only the zero vector.

Since every normal cone contains the zero vector, also all local minima/maxima meet this condition, making it indistinguishable from the min-min or max-max condition. The distinction arises at the boundary because the directions of $\nabla_\Theta(\mathbf{x}^*)$ and $\nabla_\Omega(\mathbf{x}^*)$ remain consistent in min-min or max-max problems but diverge in the min-max problem. These distinct directions make the min-max problem non-trivial. A significant drawback of local Nash equilibria, similar to Nash equilibria, is their possible nonexistence, even when dealing with well-behaved functions [10]. Additionally, establishing the existence of such points is also an NP-hard problem [2].

Transitioning to the approximate notion of a local Nash equilibrium, it's worth noting that it is defined in a very similar manner.

Definition 7 Let $\epsilon, \delta > 0$. The point $(\boldsymbol{\theta}^*, \boldsymbol{\omega}^*)$ is called an *approximate local min-max equilibrium* if it satisfies:

Also called (ϵ, δ)-local min-max equilibrium [1].

$$
\begin{aligned}
f(\boldsymbol{\theta}^*, \boldsymbol{\omega}^*) &< f(\boldsymbol{\theta}, \boldsymbol{\omega}^*) + \epsilon \quad \text{for all } \boldsymbol{\theta} \in \Theta \text{ such that } \|\boldsymbol{\theta} - \boldsymbol{\theta}^*\| \leqslant \delta, \\
f(\boldsymbol{\theta}^*, \boldsymbol{\omega}^*) &> f(\boldsymbol{\theta}^*, \boldsymbol{\omega}) - \epsilon \quad \text{for all } \boldsymbol{\omega} \in \Omega \text{ such that } \|\boldsymbol{\omega} - \boldsymbol{\omega}^*\| \leqslant \delta.
\end{aligned}
$$

The condition says that when Min-player changes their strategy within the δ-neighborhood, their loss function can increase only by ϵ. Analogous statement applies to Max-player. The advantage of this concept is that the solution is guaranteed to exist under some assumptions for f and for the locality parameter δ [1]. Specifically, the assumptions are:

1. f is G-Lipschitz, i.e., there exists $G \geqslant 0$ such that
$$|f(\mathbf{x}) - f(\mathbf{y})| \leqslant G \cdot \|\mathbf{x} - \mathbf{y}\| \quad \text{holds for all } \mathbf{x}, \mathbf{y} \in K,$$
2. f is L-smooth, i.e., there exists $L \geqslant 0$ such that
$$\|\nabla f(\mathbf{x}) - \nabla f(\mathbf{y})\| \leqslant L \cdot \|\mathbf{x} - \mathbf{y}\| \quad \text{holds for all } \mathbf{x}, \mathbf{y} \in K,$$
3. $\delta \leqslant \sqrt{\frac{2\epsilon}{L}}$.

This regime is also referred to as the local regime [1].

Approximate local min-max equilibria correspond to the approximate fixed points of the projected gradient descent-ascent dynamics (see [1, Theorem 5.1]). Let $\Pi_C \colon \mathbb{R}^m \to C$ denote the Euclidean projection onto a nonempty and compact convex set $C \subseteq \mathbb{R}^m$. In this context, we define the projected gradient descent-ascent dynamics as follows:

Definition 8 The *projected gradient descent-ascent dynamics* is defined as the map $F_{GDA} \colon K \to K$

$$F_{GDA}(\mathbf{x}) := \begin{pmatrix} \Pi_\Theta(\boldsymbol{\theta} - \nabla_\Theta f(\mathbf{x})) \\ \Pi_\Omega(\boldsymbol{\omega} + \nabla_\Omega f(\mathbf{x})) \end{pmatrix}$$

for all $\mathbf{x} = (\boldsymbol{\theta}, \boldsymbol{\omega}) \in K$.

The projection acts separately on vectors $\boldsymbol{\theta} - \nabla_\Theta f(\mathbf{x})$ and $\boldsymbol{\omega} + \nabla_\Omega f(\mathbf{x})$. It is equivalent to projecting the pair $(\boldsymbol{\theta} - \nabla_\Theta f(\mathbf{x}), \boldsymbol{\omega} + \nabla_\Omega f(\mathbf{x}))$ jointly onto K, as can be observed from the equivalence between items 4. and 5. in Proposition 7.

Due to the possible nonexistence and complexity of finding (local) Nash equilibria, researchers focus on first-order solutions. We will outline a first-order solution concept for Nash equilibrium, which has been proposed in various articles by different authors. Thus, it does not have a unified name. We will adopt the term *min-max critical point* used in [2]. Alternatively, it is also referred to as the *first-order Nash equilibrium* [3], [4], or the *game-stationary* solution [5]. The definitions in [3], [4], and [5] are identical but differ slightly from that of the min-max critical point. We will prove that the concepts are equivalent (see the equivalence between items 1. and 2. in Proposition 7).

Min-max critical points are solutions to a specific non-monotone variational inequality. The *variational inequality problem* $\mathrm{VI}(F, K)$ is to find a point $\mathbf{x}^* \in K$ such that

$$\langle F(\mathbf{x}^*), \mathbf{x}^* - \mathbf{x} \rangle \geqslant 0 \quad \text{for all } \mathbf{x} \in K,$$

where $F \colon K \to \mathbb{R}^n$ is any continuous mapping. Problem $\mathrm{VI}(F, K)$ is in general extremely difficult to solve if the mapping F fails to be monotone. In this context, we call F *monotone* if

$$\langle F(\mathbf{x}) - F(\mathbf{y}), \mathbf{x} - \mathbf{y} \rangle \geqslant 0 \qquad \text{for all } \mathbf{x}, \mathbf{y} \in K.$$

In this section we assume f to be a continuously differentiable function. Denote the set $\{1, \ldots, n\}$ by $[n]$. Each coordinate $1 \leqslant i \leqslant d_1$ and $d_1 + 1 \leqslant j \leqslant n$ is called a *minimizing* and *maximizing* coordinate, respectively. Now we can describe the min-max critical point. Define the mapping $V = (V_1, \ldots, V_n) \colon K \to \mathbb{R}^n$ by

$$V_i(\mathbf{x}) = \begin{cases} -\frac{\partial f(\mathbf{x})}{\partial x_i} & i \text{ is minimizing}, \\ \frac{\partial f(\mathbf{x})}{\partial x_i} & i \text{ is maximizing}, \end{cases} \qquad i \in [n].$$

Definition 9 A point $\mathbf{x}^* \in K$ is called a *min-max critical point* if it is a solution to the variational inequality

$$\langle V(\mathbf{x}^*), \mathbf{x}^* - \mathbf{x} \rangle \geqslant 0 \qquad \text{for all } \mathbf{x} \in K. \qquad \mathrm{VI}(V, K)$$

We present a simple example to demonstrate the nature of the min-max critical points.

Example 1 Let $f(\theta, \omega) = (\theta - \omega)^2$, $(\theta, \omega) \in [0, 1]^2$. This function has no local Nash equilibria. Min-max critical points of f are all points on the diagonal connecting the points $(0, 0)$ and $(1, 1)$.

When the function f is (strongly) convex in $\boldsymbol{\theta}$ and (strongly) concave in $\boldsymbol{\omega}$, the mapping V is (strongly) monotone [2], therefore, classical methods for solving variational inequalities can be applied [12].

2 A strongly monotone mapping $F(\cdot)$ satisfies $\langle F(\mathbf{x}) - F(\mathbf{y}), \mathbf{x} - \mathbf{y} \rangle \geqslant \sigma \|\mathbf{x} - \mathbf{y}\|^2$ for all $\mathbf{x}, \mathbf{y} \in K$, where $\sigma > 0$. If it satisfies this inequality for $\sigma = 0$, the corresponding VI is monotone.

However, in the general case when f may fail to be convex-concave, the monotonicity property of V can no longer hold. The continuity of V ensures the existence of a solution to $\mathrm{VI}(V, K)$, as it corresponds to the fixed point of the projected gradient descent-ascent dynamics (see the equivalence between items 1. and 4. in Proposition 7). The existence of this fixed point is guaranteed by Brouwer's fixed point theorem. Under the assumptions for f stated below Definition 7, the solution to $\mathrm{VI}(V, K)$ also represents an approximate local min-max equilibrium whenever $\delta \leqslant \sqrt{\frac{2\epsilon}{L}}$. The complexity of computing min-max critical points is equivalent to Brouwer's fixed points computation, which lies in the PPAD [6] class. However, verifying if a point is min-max critical is easy and can be done in linear time.

Let us outline the definition of the first-order Nash equilibrium.

Definition 10 The *first-order Nash equilibrium (FNE)* is a point $\mathbf{x}^* = (\boldsymbol{\theta}^*, \boldsymbol{\omega}^*) \in K$ which satisfies the following conditions:

$$\langle \nabla_\Theta f(\mathbf{x}^*), \boldsymbol{\theta} - \boldsymbol{\theta}^* \rangle \geqslant 0, \quad \forall \boldsymbol{\theta} \in \Theta,$$

$$\langle \nabla_\Omega f(\mathbf{x}^*), \boldsymbol{\omega} - \boldsymbol{\omega}^* \rangle \leqslant 0, \quad \forall \boldsymbol{\omega} \in \Omega.$$

One can verify that this is equivalent to the first-order necessary condition for a local NE (Proposition 6). For $\mathbf{x} \in K$ and $i \in [n]$, define the set $K_i(\mathbf{x}) = \{x_i' \in \mathbb{R} \mid (x_1, \ldots, x_{i-1}, x_i', x_{i+1}, \ldots, x_n) \in K\}$. We consolidate the properties of the min-max critical point in the following proposition.

Proposition 7 Let $\mathbf{x}^* = (\boldsymbol{\theta}^*, \boldsymbol{\omega}^*) \in K$. The following are equivalent.

1. The point $\mathbf{x}^*$ is min-max critical.

2. The point $\mathbf{x}^*$ is a first-order Nash equilibrium.

3. For each $i \in [n]$ and every $x_i \in K_i$, $V_i(\mathbf{x}^*)(x_i^* - x_i) \geqslant 0$.

4. $\mathbf{x}^* = \Pi_K(\mathbf{x}^* + V(\mathbf{x}^*))$.

5. $\boldsymbol{\theta}^* = \Pi_\Theta(\boldsymbol{\theta}^* - \nabla_\Theta f(\mathbf{x}^*))$ and $\boldsymbol{\omega}^* = \Pi_\Omega(\boldsymbol{\omega}^* + \nabla_\Omega f(\mathbf{x}^*))$, i.e., $\mathbf{x}^* = F_{GDA}(\mathbf{x}^*)$

Proof 1. $\Rightarrow$ 2. Consider an arbitrary $\boldsymbol{\theta} \in \Theta$. Choose $\mathbf{x} = (\boldsymbol{\theta}, \boldsymbol{\omega}^*) \in K$. Then

$$\langle -\nabla_\Theta f(\mathbf{x}^*), \boldsymbol{\theta}^* - \boldsymbol{\theta} \rangle = \langle V(\mathbf{x}^*), \mathbf{x}^* - \mathbf{x} \rangle \geqslant 0,$$

which is equivalent with $\langle \nabla_\Theta f(\mathbf{x}^*), \boldsymbol{\theta} - \boldsymbol{\theta}^* \rangle \geqslant 0$. Consider an arbitrary $\boldsymbol{\omega} \in \Omega$. Choose $\mathbf{x} = (\boldsymbol{\theta}^*, \boldsymbol{\omega}) \in K$. Then

$$\langle \nabla_\Omega f(\mathbf{x}^*), \boldsymbol{\omega}^* - \boldsymbol{\omega} \rangle = \langle V(\mathbf{x}^*), \mathbf{x}^* - \mathbf{x} \rangle \geqslant 0,$$

which is equivalent with $\langle \nabla_\Omega f(\mathbf{x}^*), \boldsymbol{\omega} - \boldsymbol{\omega}^* \rangle \leqslant 0$.

2. $\Rightarrow$ 1. Let $\mathbf{x} = (\boldsymbol{\theta}, \boldsymbol{\omega}) \in K$. We know that $\langle -\nabla_\Theta f(\mathbf{x}^*), \boldsymbol{\theta}^* - \boldsymbol{\theta} \rangle \geqslant 0$ and $\langle \nabla_\Omega f(\mathbf{x}^*), \boldsymbol{\omega}^* - \boldsymbol{\omega} \rangle \geqslant 0$. After summing up the two inequalities we get $\langle V(\mathbf{x}^*), \mathbf{x}^* - \mathbf{x} \rangle \geqslant 0$.

1. $\Rightarrow$ 3. Let $i \in [n]$ and $x_i \in K_i$. Define $\mathbf{x} \in K$ by $x_j = x_j^*$ for all $j \neq i$ and $x_j = x_i$ when $j = i$. Then

$$V_i(\mathbf{x}^*)(x_i^* - x_i) = \langle V(\mathbf{x}^*), \mathbf{x}^* - \mathbf{x} \rangle \geqslant 0.$$

3. $\Rightarrow$ 1. This follows directly from the assumption,

$$\langle V(\mathbf{x}^*), \mathbf{x}^* - \mathbf{x} \rangle = \sum_{i=1}^{n} V_i(\mathbf{x}^*)(x_i^* - x_i) \geqslant 0.$$

1. $\Leftrightarrow$ 4. The wellknown characterization of projection says that 4. holds true if, and only if, the inequality

$$\langle \mathbf{x}^* + V(\mathbf{x}^*) - \mathbf{x}^*, \mathbf{x} - \mathbf{x}^* \rangle \leqslant 0$$

is valid for every $\mathbf{x} \in K$, which is equivalent with the definition of the min-max critical point $\mathbf{x}^*$.

4. $\Leftrightarrow$ 5. By the definition of the Euclidean projection,

$$\Pi_{\Theta \times \Omega} \begin{pmatrix} \boldsymbol{\theta}^* - \nabla_\Theta f(\mathbf{x}^*) \\ \boldsymbol{\omega}^* + \nabla_\Omega f(\mathbf{x}^*) \end{pmatrix} = \underset{(\boldsymbol{\theta}, \boldsymbol{\omega}) \in \Theta \times \Omega}{\operatorname{argmin}} \left\| \begin{pmatrix} \boldsymbol{\theta}^* - \nabla_\Theta f(\mathbf{x}^*) - \boldsymbol{\theta} \\ \boldsymbol{\omega}^* + \nabla_\Omega f(\mathbf{x}^*) - \boldsymbol{\omega} \end{pmatrix} \right\|^2$$

$$= \underset{(\boldsymbol{\theta}, \boldsymbol{\omega}) \in \Theta \times \Omega}{\operatorname{argmin}} \left(\sum_{i=1}^{d_1} (\theta_i^* - \nabla_\Theta f(\mathbf{x}^*)_i - \theta_i)^2 + \sum_{i=d_1+1}^{d_1+d_2} (\omega_i + \nabla_\Omega f(\mathbf{x}^*)_i - \omega_i)^2 \right)$$

$$= \underset{(\boldsymbol{\theta}, \boldsymbol{\omega}) \in \Theta \times \Omega}{\operatorname{argmin}} \left(\|\boldsymbol{\theta}^* - \nabla_\Theta f(\mathbf{x}^*) - \boldsymbol{\theta}\|^2 + \|\boldsymbol{\omega}^* + \nabla_\Omega f(\mathbf{x}^*) - \boldsymbol{\omega}\|^2 \right).$$

Since the two summands do not contain $\boldsymbol{\theta}$, $\boldsymbol{\omega}$ simultaneously, it is equal to

$$\begin{pmatrix} \underset{\boldsymbol{\theta} \in \Theta}{\operatorname{argmin}} \|\boldsymbol{\theta}^* - \nabla_\Theta f(\mathbf{x}^*) - \boldsymbol{\theta}\|^2 \\ \underset{\boldsymbol{\omega} \in \Omega}{\operatorname{argmin}} \|\boldsymbol{\omega}^* + \nabla_\Omega f(\mathbf{x}^*) - \boldsymbol{\omega}\|^2 \end{pmatrix} = \begin{pmatrix} \Pi_\Theta(\boldsymbol{\theta}^* - \nabla_\Theta f(\mathbf{x}^*)) \\ \Pi_\Omega(\boldsymbol{\omega}^* + \nabla_\Omega f(\mathbf{x}^*)) \end{pmatrix}.$$

The equivalence between items 1. and 4. indicates that the min-max critical point corresponds to the fixed point of the projected gradient descent-ascent dynamics with a step size 1. This equivalence also holds for a step size $\alpha > 0$, which can be proved by using a similar argument as in the proof above.

We will use a particular description of the min-max critical points in the case of a unit hypercube $K = [0, 1]^n$ [2].

Definition 11 Let $\mathbf{x} \in [0,1]^n$. The coordinate $i \in [n]$ is *satisfied* at $\mathbf{x}$ if one of the following conditions hold:

1. $V_i(\mathbf{x}) = 0$.
2. $V_i(\mathbf{x}) < 0$ and $x_i = 0$.
3. $V_i(\mathbf{x}) > 0$ and $x_i = 1$.

In case that item 1. is satisfied, then i is called *zero-satisfied* at $\mathbf{x}$. If 2. or 3. are true, then i is called *boundary-satisfied* at $\mathbf{x}$.

In Definition 11 and Proposition 8, we consider $K = [0,1]^n$ because the method STON'R, presented in Chapter 3, operates on the unit hypercube.

Proposition 8 Let $\mathbf{x}^* \in [0,1]^n$. The following are equivalent.

1. The point $\mathbf{x}^*$ is min-max critical.
2. Each coordinate $i \in [n]$ is satisfied at $\mathbf{x}^*$.

Proposition 8 is the consequence of the equivalence between items 1. and 3. of Proposition 7.

Let K be a compact convex set. We will need the approximate counterparts of the concepts introduced above.

Definition 12 We say that a point $\mathbf{x} \in K$ is an *α-approximate solution to* VI(V, K) for some $\alpha > 0$ if

$$\langle V(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle \leqslant \alpha \quad \text{for all } \mathbf{y} \in K.$$

Clearly, for every $\alpha > 0$ any min-max critical point $\mathbf{x}$ is an α-approximate solution to VI(V, K).

Definition 13 We call a point $\mathbf{x} \in K$ an *α-approximate fixed point* of the projected gradient descent-ascent dynamics for some $\alpha > 0$ if $\|F_{GDA}(\mathbf{x}) - \mathbf{x}\| \leqslant \alpha$.

Proposition 9 If $\mathbf{x} \in K$ is an α-approximate solution to VI(V, K), then it is an $\sqrt{\alpha}$-approximate fixed point of the projected gradient descent-ascent dynamics.

Proof Let $\mathbf{x} \in K$ be an α-approximate solution to VI(V, K) and let $\mathbf{z} \in K$ be defined as $\mathbf{z} = \Pi_K(\mathbf{x} + V(\mathbf{x}))$. We have

$$\begin{aligned}
\alpha &\geqslant \langle V(\mathbf{x}), \mathbf{z} - \mathbf{x} \rangle = \langle \mathbf{z} - \mathbf{x} + \mathbf{x} + V(\mathbf{x}) - \mathbf{z}, \mathbf{z} - \mathbf{x} \rangle \\
&= \|\mathbf{z} - \mathbf{x}\|^2 - \langle \mathbf{x} + V(\mathbf{x}) - \Pi_K(\mathbf{x} + V(\mathbf{x})), \mathbf{x} - \Pi_K(\mathbf{x} + V(\mathbf{x})) \rangle \\
&\geqslant \|\mathbf{z} - \mathbf{x}\|^2 = \|\Pi_K(\mathbf{x} + V(\mathbf{x})) - \mathbf{x}\|^2.
\end{aligned}$$

The last inequality follows from the property of the Euclidean projection. Therefore, $\|\Pi_K(\mathbf{x} + V(\mathbf{x})) - \mathbf{x}\| \leqslant \sqrt{\alpha}$.

The converse proposition also stands: an α-approximate fixed point is a $(c\alpha)$-approximate solution to VI(V, K), where the constant c depends on both the diameter of the strategy set and the maximum norm of V [13, Proposition 3.1]. We will utilize Proposition 9 in the discrete dynamics of the STay-ON-the-Ridge algorithm.

In this section, we will extend the concept of the min-max critical point to locally Lipschitz functions. These functions are continuous, and by Rademacher's theorem, they are also differentiable almost everywhere. That is, the set of points where the function is not differentiable has Lebesgue measure zero. We will need the tools of *Clarke nonsmooth analysis* [14].

Let us start by introducing the Clarke generalized derivative and gradient [14]. Let $f\colon \mathbb{R}^n \to \mathbb{R}$ be Lipschitz near a given point $\mathbf{x} \in \mathbb{R}^n$, that is, for some $l, \epsilon \geqslant 0$, we have

$$|f(\mathbf{y}) - f(\mathbf{z})| \leqslant l\|\mathbf{y} - \mathbf{z}\|$$

for all $\mathbf{y}, \mathbf{z}$ in $\mathbb{R}^n$ such that $\|\mathbf{y} - \mathbf{x}\| \leqslant \epsilon$ and $\|\mathbf{z} - \mathbf{x}\| \leqslant \epsilon$. The *Clarke generalized directional derivative* of f at $\mathbf{x}$ in the direction $\mathbf{h}$, denoted as $f^\circ(\mathbf{x}, \mathbf{h})$, is defined as follows:

$$f^\circ(\mathbf{x}, \mathbf{h}) = \limsup_{\substack{\mathbf{x}' \to \mathbf{x} \\ t \downarrow 0}} \frac{f(\mathbf{x}' + t \cdot \mathbf{h}) - f(\mathbf{x}')}{t},$$

where $\mathbf{x}' \in \mathbb{R}^n$ and $t > 0$. Then, using the above definition of f°, the *Clarke generalized gradient* $\partial_C f(\mathbf{x})$ of f at $\mathbf{x}$ (also called *Clarke subdifferential*) is given as

$$\partial_C f(\mathbf{x}) = \{\mathbf{y} \in \mathbb{R}^n \mid \langle \mathbf{y}, \mathbf{h} \rangle \leqslant f^\circ(\mathbf{x}, \mathbf{h}),\ \forall \mathbf{h} \in \mathbb{R}^n\}.$$

We have therefore

$$f^\circ(\mathbf{x}, \mathbf{h}) = \max\{\langle \mathbf{y}, \mathbf{h} \rangle \mid \mathbf{y} \in \partial_c f(\mathbf{x})\} \quad \forall \mathbf{h} \in \mathbb{R}^n.$$

The Clarke generalized gradient is a set-valued function (multifunction). The following theorem [14] states that the Clarke generalized gradient $\partial_C f(\mathbf{x})$ can be derived from the values of $\nabla f(\mathbf{u})$ at nearby points $\mathbf{u}$ where $f'(\mathbf{u})$ exists. Moreover, the construction remains unaffected by points $\mathbf{u}$ belonging to any set of measure zero.

Theorem 1 (Gradient formula) Let $\mathbf{x} \in \mathbb{R}^n$ and $f\colon \mathbb{R}^n \to \mathbb{R}$ be Lipschitz near $\mathbf{x}$. Let E be any subset of zero measure in $\mathbb{R}^n$, and let E_f be the set of points at which f fails to be differentiable. Then

$$\partial_C f(\mathbf{x}) = \mathrm{co}\left\{\lim_{i \to \infty} \nabla f(\mathbf{x}_i) \mid \mathbf{x}_i \to \mathbf{x},\ \mathbf{x}_i \notin E \cup E_f\right\},$$

where co denotes the convex hull.

We also need to outline generalized variational inequalities, drawing upon the formulation provided in [15]. Consider a set-valued function γ from some subset C of $\mathbb{R}^n$ to the family of subsets of $\mathbb{R}^n$.

The *Generalized variational inequality problem* GVI(γ, C) is to calculate a solution $(\mathbf{x}, \mathbf{y})$ satisfying the following:

$$\begin{aligned} &1.\quad \mathbf{x} \in C, \\ &2.\quad \langle \mathbf{y}, \mathbf{x} - \mathbf{x}' \rangle \geqslant 0, \text{ for each } \mathbf{x}' \in C, \\ &3.\quad \mathbf{y} \in \gamma(\mathbf{x}). \end{aligned} \tag{GVI(γ, C)}$$

Next, we introduce the definitions of contractibility and upper-semicontinuity. A subset S of $\mathbb{R}^n$ is *contractible*, if there is an $\mathbf{x}^\circ \in S$ and a continuous function $g \colon S \times [0,1] \to S$, such that

$$g(\mathbf{x}, 0) = \mathbf{x} \text{ and } g(\mathbf{x}, 1) = \mathbf{x}^\circ, \text{ for each } \mathbf{x} \in S.$$

If S is convex, then S is contractible, since $g(\mathbf{x}, t) = (1-t)\mathbf{x} + t\mathbf{x}^\circ$ has the described property for any $\mathbf{x}^\circ \in S$. A multifunction γ on $C \subseteq \mathbb{R}^n$ is *upper-semicontinuous* if the following holds: if a sequence of vectors $\mathbf{x}^k$ in C converges to $\mathbf{x} \in C$ and a sequence of vectors $\mathbf{y}^k \in \gamma(\mathbf{x}^k)$ converges to $\mathbf{y}$, then $\mathbf{y} \in \gamma(\mathbf{x})$. We are now prepared to state the existence theorem.

Theorem 2 (Hartman-Stampacchia, Saigal) Assume that

1. C is a nonempty, compact and convex set in $\mathbb{R}^n$,
2. γ is an upper-semicontinuous mapping from C to the family of subsets of $\mathbb{R}^n$,
3. $\gamma(\mathbf{x})$ is a nonempty, compact, and contractible set in $\mathbb{R}^n$ for each $\mathbf{x} \in C$.

Then there is a solution to the variational inequality GVI(γ, C).

We establish a generalized version of the mapping V introduced in Section 2.3. Recall that K is a convex compact set. Define a mapping $\mathcal{V} \colon K \twoheadrightarrow \mathbb{R}^n$ as the composition of the Clarke subdifferential and a linear function $L = (L_1, \dots, L_n) \colon \mathbb{R}^n \to \mathbb{R}^n$ given by

$$L_i(\mathbf{y}) = \begin{cases} -y_i & i \text{ is minimizing}, \\ y_i & i \text{ is maximizing}, \end{cases} \qquad i \in [n].$$

Let $\mathbf{x} \in K$. Define $\mathcal{V}(\mathbf{x})$ as $\bigcup\limits_{\mathbf{y} \in \partial_C f(\mathbf{x})} \{L(\mathbf{y})\}$.

The mapping $\mathcal{V}$ is indeed a generalization of V. In the case when f is continuously differentiable, the Clarke generalized gradient reduces to a single point, representing the standard gradient. Now, we can define the generalized version of the min-max critical point.

Definition 14 A point $\mathbf{x}^* \in K$ is called a *generalized min-max critical point* if there exists a point $\mathbf{y}^* \in \mathcal{V}(\mathbf{x}^*)$ such that

$$\langle \mathbf{y}^*, \mathbf{x}^* - \mathbf{x} \rangle \geqslant 0 \text{ for each } \mathbf{x} \in K. \tag{GVI$(\mathcal{V}, K)$}$$

Similar to the smooth min-max critical point, within the interior of K, the point $\mathbf{y}^*$ must be the zero vector to satisfy condition $\mathrm{GVI}(\mathcal{V}, K)$.

Now we show that the $\mathrm{GVI}(\mathcal{V}, K)$ has a solution.

Proposition 10 There exists a generalized min-max critical point, i.e., there exists a solution to $\mathrm{GVI}(\mathcal{V}, K)$.

Proof Utilizing Theorem 2, we establish the existence of the generalized min-max critical point by setting $C := K$ and $\gamma := \mathcal{V}$. This leads us to assert that each of the conditions stated in items 1-3 is indeed fulfilled.

1. The condition holds directly from the assumptions for K.

2. The Clarke generalized gradient ∂_C is an upper-semicontinuous set-valued function [16]. The linear function L is continuous, and therefore also upper-semicontinuous. Since the mapping $\mathcal{V}$ is the composition of ∂_C and L, it is also upper-semicontinuous [17, Theorem 17.23].

3. Considering an arbitrary point $\mathbf{x} \in K$, the set $\mathcal{V}(\mathbf{x})$ is nonempty due to the nonemptiness of $\partial_C f(\mathbf{x})$ [14, Proof of Theorem 10.27]. The compactness of $\mathcal{V}(\mathbf{x})$ is derived from the compactness of $\partial_C f(\mathbf{x})$ and the continuity of L, a consequence of its linearity. This follows directly from a well-known theorem in analysis stating that the continuous image of a compact set remains compact. To show contractibility, we use the property that every convex set contracts. We now establish convexity of $\mathcal{V}(\mathbf{x})$ for an arbitraty $\mathbf{x} \in K$. Consider arbitrary points $\mathbf{r}'$ and $\mathbf{s}'$ in $\mathcal{V}(\mathbf{x})$. Then there exist $\mathbf{r}$ and $\mathbf{s}$ in $\partial_C f(\mathbf{x})$ such that $\mathbf{r}' = L(\mathbf{r})$ and $\mathbf{s}' = L(\mathbf{s})$. Let $\lambda \in [0, 1]$. We want to prove that $\lambda \mathbf{r}' + (1 - \lambda)\mathbf{s}' \in \mathcal{V}(\mathbf{x})$. Since $\partial_C f(\mathbf{x})$ is convex, it holds $\lambda \mathbf{r} + (1 - \lambda)\mathbf{s} \in \partial_C f(\mathbf{x})$. From the linearity of L we get $L(\lambda \mathbf{r} + (1 - \lambda)\mathbf{s}) = \lambda L(\mathbf{r}) + (1 - \lambda)L(\mathbf{s}) = \lambda \mathbf{r}' + (1 - \lambda)\mathbf{s}' \in \mathcal{V}(\mathbf{x})$.

Now we can proceed to define and investigate analogous properties that hold for the smooth min-max critical point. Let $\mathbf{x}_{\mathrm{min}}$ and $\mathbf{x}_{\mathrm{max}}$ be the restrictions of a vector $\mathbf{x} \in \mathbb{R}^n$ to the minimizing and maximizing coordinates, respectively.

Definition 15 A point $\mathbf{x}^* = (\boldsymbol{\theta}^*, \boldsymbol{\omega}^*) \in K$ is called a *generalized first-order Nash equilibrium (GFNE)* if there exists a point $\mathbf{y}^* \in \partial_C f(\mathbf{x}^*)$ such that

$$\langle \mathbf{y}^*_{\mathrm{min}}, \boldsymbol{\theta} - \boldsymbol{\theta}^* \rangle \geqslant 0, \quad \forall \boldsymbol{\theta} \in \Theta,$$

$$\langle \mathbf{y}^*_{\mathrm{max}}, \boldsymbol{\omega} - \boldsymbol{\omega}^* \rangle \leqslant 0, \quad \forall \boldsymbol{\omega} \in \Omega.$$

In this context we can establish a first-order necessary condition for a local Nash equilibrium (Definition 5) for nondifferentiable functions.

Proposition 11 (First-order necessary condition) Any local Nash equilibrium $\mathbf{x}^* = (\boldsymbol{\theta}^*, \boldsymbol{\omega}^*) \in K$ meets $-\mathbf{y}^*_{\min} \in N_\Theta(\boldsymbol{\theta}^*)$ and $\mathbf{y}^*_{\max} \in N_\Omega(\boldsymbol{\omega}^*)$ for some $\mathbf{y}^* \in \partial_C f(\mathbf{x}^*)$, where $N_\Theta(\boldsymbol{\theta}^*)$ and $N_\Omega(\boldsymbol{\omega}^*)$ are normal cones (see Definition 6).

Proposition 11 is derived from [18]. Analogous to the first-order Nash equilibrium for differentiable functions (Definition 10), the condition in Proposition 11 is equivalent to the generalized first-order Nash equilibrium.

Proposition 12 Let $\mathbf{x}^* = (\boldsymbol{\theta}^*, \boldsymbol{\omega}^*) \in K$. The following are equivalent.

1. The point $\mathbf{x}^*$ is generalized min-max critical.
2. The point $\mathbf{x}^*$ is a generalized first-order Nash equilibrium.
3. For each $i \in [n]$ and every $x_i \in K_i$, $y^*_i(x^*_i - x_i) \geqslant 0$, for some $\mathbf{y}^* \in \mathcal{V}(\mathbf{x}^*)$.
4. $\mathbf{x}^* = \Pi_K(\mathbf{x}^* + \mathbf{y}^*)$, for some $\mathbf{y}^* \in \mathcal{V}(\mathbf{x}^*)$.
5. $\boldsymbol{\theta}^* = \Pi_\Theta(\boldsymbol{\theta}^* + \mathbf{y}^*_{\min})$ and $\boldsymbol{\omega}^* = \Pi_\Omega(\boldsymbol{\omega}^* + \mathbf{y}^*_{\max})$, for some $\mathbf{y}^* \in \mathcal{V}(\mathbf{x}^*)$.

Proof The proof follows a similar line of reasoning in the proof of Proposition 7.

1. $\Rightarrow$ 2. There exists some $\mathbf{y}^* \in \mathcal{V}(\mathbf{x}^*)$ such that $\langle \mathbf{y}^*, \mathbf{x}^* - \mathbf{x}\rangle \geqslant 0$ for all $\mathbf{x} \in K$. Let $\mathbf{z}^* = (-\mathbf{y}^*_{\min}, \mathbf{y}^*_{\max}) \in \partial_C f(\mathbf{x}^*)$. Consider an arbitraty $\boldsymbol{\theta} \in \Theta$. Choose $\mathbf{x} = (\boldsymbol{\theta}, \boldsymbol{\omega}^*) \in K$. Then

$$\langle \mathbf{y}^*_{\min}, \boldsymbol{\theta}^* - \boldsymbol{\theta}\rangle = \langle \mathbf{y}^*, \mathbf{x}^* - \mathbf{x}\rangle \geqslant 0,$$

which is equivalent with $\langle \mathbf{z}^*_{\min}, \boldsymbol{\theta} - \boldsymbol{\theta}^*\rangle \geqslant 0$. Consider an arbitrary $\boldsymbol{\omega} \in \Omega$. Choose $\mathbf{x} = (\boldsymbol{\theta}^*, \boldsymbol{\omega}) \in K$. Then

$$\langle \mathbf{y}^*_{\max}, \boldsymbol{\omega}^* - \boldsymbol{\omega}\rangle = \langle \mathbf{y}^*, \mathbf{x}^* - \mathbf{x}\rangle \geqslant 0,$$

which is equivalent with $\langle \mathbf{z}^*_{\max}, \boldsymbol{\omega} - \boldsymbol{\omega}^*\rangle \leqslant 0$.

2. $\Rightarrow$ 1. Let $\mathbf{x} = (\boldsymbol{\theta}, \boldsymbol{\omega}) \in K$. There exists some $\mathbf{z}^* \in \partial_C f(\mathbf{x}^*)$ such that $\langle -\mathbf{z}^*_{\min}, \boldsymbol{\theta}^* - \boldsymbol{\theta}\rangle \geqslant 0$ and $\langle \mathbf{z}^*_{\max}, \boldsymbol{\omega}^* - \boldsymbol{\omega}\rangle \geqslant 0$. After summing up the two inequalities we get $\langle \mathbf{y}^*, \mathbf{x}^* - \mathbf{x}\rangle \geqslant 0$, where $\mathbf{y}^* = (-\mathbf{z}^*_{\min}, \mathbf{z}^*_{\max}) \in \mathcal{V}(\mathbf{x}^*)$.

1. $\Rightarrow$ 3. Since the point $\mathbf{x}^*$ is generalized min-max critical, there exist some $\mathbf{y}^* \in \mathcal{V}(\mathbf{x}^*)$ such that $(\mathbf{x}^*, \mathbf{y}^*)$ is a solution to the GVI$(\mathcal{V}, K)$. Let $i \in [n]$ and $x_i \in K_i$. Define $\mathbf{x} \in K$ by $x_j = x^*_j$ for all $j \neq i$ and $x_j = x_i$ when $j = i$. Then

$$y^*_i(x^*_i - x_i) = \langle \mathbf{y}^*, \mathbf{x}^* - \mathbf{x}\rangle \geqslant 0.$$

3. $\Rightarrow$ 1. This follows directly from the assumption,

$$\langle \mathbf{y}^*, \mathbf{x}^* - \mathbf{x}\rangle = \sum_{i=1}^{n} y^*_i(x^*_i - x_i) \geqslant 0.$$

1. ⇔ 4. The wellknown characterization of projection says that 4. holds true if, and only if, the inequality

$$\langle \mathbf{x}^* + \mathbf{y}^* - \mathbf{x}^*, \mathbf{x} - \mathbf{x}^* \rangle \leqslant 0$$

is valid for every $\mathbf{x} \in K$, which is equivalent with the definition of the generalized min-max critical point $\mathbf{x}^*$.

4. ⇔ 5. By the definition of the Euclidean projection,

$$\Pi_{\Theta \times \Omega} \begin{pmatrix} \boldsymbol{\theta}^* + \mathbf{y}^*_{\min} \\ \boldsymbol{\omega}^* + \mathbf{y}^*_{\max} \end{pmatrix} = \operatorname*{argmin}_{(\boldsymbol{\theta}, \boldsymbol{\omega}) \in \Theta \times \Omega} \left\| \begin{pmatrix} \boldsymbol{\theta}^* + \mathbf{y}^*_{\min} - \boldsymbol{\theta} \\ \boldsymbol{\omega}^* + \mathbf{y}^*_{\max} - \boldsymbol{\omega} \end{pmatrix} \right\|^2$$

$$= \operatorname*{argmin}_{(\boldsymbol{\theta}, \boldsymbol{\omega}) \in \Theta \times \Omega} \left(\sum_{i=1}^{d_1} (\theta_i^* + y_i^* - \theta_i)^2 + \sum_{i=d_1+1}^{d_1+d_2} (\omega_i^* + y_i^* - \omega_i)^2 \right)$$
$$= \operatorname*{argmin}_{(\boldsymbol{\theta}, \boldsymbol{\omega}) \in \Theta \times \Omega} \left(\|\boldsymbol{\theta}^* + \mathbf{y}^*_{\min} - \boldsymbol{\theta}\|^2 + \|\boldsymbol{\omega}^* + \mathbf{y}^*_{\max} - \boldsymbol{\omega}\|^2 \right).$$

Since the two summands do not contain $\boldsymbol{\theta}$, $\boldsymbol{\omega}$ simultaneously, it is equal to

$$\begin{pmatrix} \operatorname*{argmin}_{\boldsymbol{\theta} \in \Theta} \|\boldsymbol{\theta}^* + \mathbf{y}^*_{\min} - \boldsymbol{\theta}\|^2 \\ \operatorname*{argmin}_{\boldsymbol{\omega} \in \Omega} \|\boldsymbol{\omega}^* + \mathbf{y}^*_{\max} - \boldsymbol{\omega}\|^2 \end{pmatrix} = \begin{pmatrix} \Pi_\Theta(\boldsymbol{\theta}^* + \mathbf{y}^*_{\min}) \\ \Pi_\Omega(\boldsymbol{\omega}^* + \mathbf{y}^*_{\max}) \end{pmatrix}.$$

The equivalence between items 1. and 4. imply that we can solve $\mathrm{GVI}(\mathcal{V}, K)$ by computing fixed points of the mapping $\phi \colon K \twoheadrightarrow \mathbb{R}^n$, where $\phi(\mathbf{x}) = \{\Pi_K(\mathbf{x} + \mathbf{y}) \mid \mathbf{y} \in \mathcal{V}(\mathbf{x})\}$. A point $\mathbf{x} \in K$ is considered a fixed point of ϕ if $\mathbf{x} \in \phi(\mathbf{x})$. This equivalence also extends to cases with a positive step size $\alpha > 0$, confirmed through a similar reasoning as in the proof above.

The following example shows that the generalized min-max critical points capture the saddle points of nonsmooth fucntions.

Example 2 Consider a two-variable piecewise affine function over the unit square $K = [0, 1]^n$ from [19, Example 3.2]. We note that the first coordinate is maximizing and the second minimizing. The function is defined as follows:

$$f(\mathbf{x}) = \begin{cases} 2x_1 - x_2 & \text{for } \mathbf{x} \in A = \{\mathbf{x} \mid x_2 \geqslant x_1 \text{ and } x_1 \leqslant \frac{1}{2}\} \\ -2x_1 - x_2 + 2 & \text{for } \mathbf{x} \in B = \{\mathbf{x} \mid x_2 \leqslant -x_1 + 1 \text{ and } x_1 \geqslant \frac{1}{2}\} \\ 1 - x_1 & \text{for } \mathbf{x} \in C = \{\mathbf{x} \mid x_2 \geqslant -x_1 + 1 \text{ and } x_2 \leqslant x_1\} \\ -2x_1 + x_2 + 1 & \text{for } \mathbf{x} \in D = \{\mathbf{x} \mid x_2 \geqslant x_1 \text{ and } x_1 \geqslant 1/2\} \\ 2x_1 + x_2 - 1 & \text{for } \mathbf{x} \in E = \{\mathbf{x} \mid x_2 \geqslant -x_1 + 1 \text{ and } x_1 \leqslant 1/2\} \\ x_1 & \text{for } \mathbf{x} \in F = \{\mathbf{x} \mid x_2 \geqslant x_1 \text{ and } x_2 \leqslant -x_1 + 1\}. \end{cases}$$

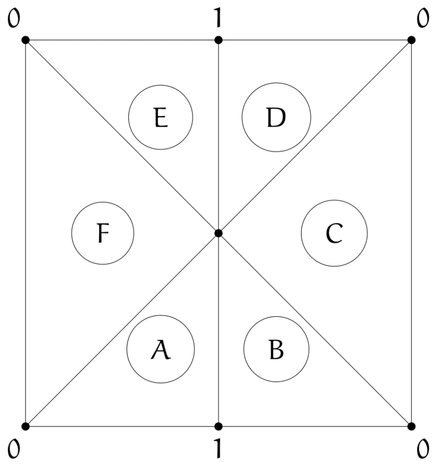

Figure 2.1: Regions of f

The Nash equilibrium of f is $\left(\frac{1}{2}, \frac{1}{2}\right)$, which can be seen from the graph below.

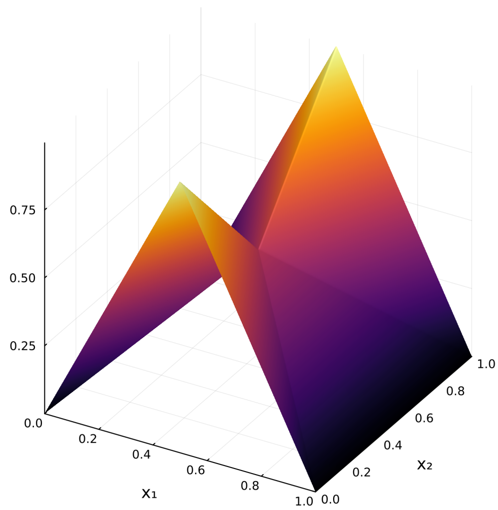

Figure 2.2: $f(\mathbf{x})$

The function meets the Lipschitz condition as it is piecewise affine. The boundaries of sets A, B, C, D, E, and F together form a set of measure 0. If $\mathbf{x}$ does not lie in this set, f is differentiable at $\mathbf{x}$, and $\nabla f(\mathbf{x})$ is one of the points $(2, -1), (-2, -1), (-1, 0), (-2, 1), (2, 1)$, or $(1, 0)$. The gradient formula implies that $\partial_C f(\frac{1}{2}, \frac{1}{2})$ forms a rectangle with vertices $(-2, -1), (2, -1), (2, 1), (-2, 1)$ obtained as a convex hull of these four points. Since this rectangle contains the origin $(0, 0)$, the point $(\frac{1}{2}, \frac{1}{2})$ is indeed a solution to the $\text{GVI}(\mathcal{V}, K)$. One can verify that this is the only generalized min-max critical point of f.

3 STAY-ON-THE-RIDGE

We describe a second-order method, called STay-ON-the-Ridge, proposed in [2]. This algorithm is guaranteed to converge to the min-max critical point (Definition 9) given that the corresponding strategy sets are unit hypercubes. STON'R requires the mapping V (Definition 9) to be L-Lipschitz and Λ-smooth. This means

$$\|V(\mathbf{x}) - V(\mathbf{y})\| \leqslant L\|\mathbf{x}-\mathbf{y}\|, \text{ for all } \mathbf{x}, \mathbf{y} \in [0,1]^n, \text{ and} \tag{2}$$

$$\|J(\mathbf{x}) - J(\mathbf{y})\|_F \leqslant \Lambda\|\mathbf{x}-\mathbf{y}\|, \text{ for all } \mathbf{x}, \mathbf{y} \in [0,1]^n, \tag{3}$$

where J is the Jacobian matrix of V, and $\|A\|_F$ denotes the Frobenius norm of the matrix A.

3.1 CONTINUOUS DYNAMICS

The objective of the algorithm is to find a point $\mathbf{x} \in [0,1]^n$ such that each coordinate $i \in [n]$ is satisfied at $\mathbf{x}$ according to Definition 11. It is initialized at $\mathbf{x}(0) = (0, \dots, 0)$ and aims to satisfy all coordinates one-by-one.

The algorithm starts *epoch* (i, S) at point $\mathbf{x}$ if all coordinates in the set $S \subseteq [i-1]$ are zero-satisfied, all coordinates in $[i-1] \setminus S$ are boundary-satisfied, and the algorithm's goal is to find a point $\mathbf{x}' \in [0,1]^n$ that satisfies all coordinates $\leqslant i$. Assume that the algorithm starts epoch (i, S) at point $\mathbf{x}(t)$ at time t. It tries to achieve the goal of epoch (i, S) starting at $\mathbf{x}(t)$ as follows:

- It tries to find a point inside the connected subset $S^i(\mathbf{x}(t)) \subseteq [0,1]^n$ that contains all points $\mathbf{z}$ s.t. (a) all coordinates in S are zero-satisfied and all coordinates in $[i-1] \setminus S$ are boundary-satisfied and (b) for all $j \geqslant i+1$, $z_j = x_j(t)$.

- It navigates $S^i(\mathbf{x}(t))$ in the hopes of satisfying also the coordinate i. It runs a continuous time dynamics $\{\mathbf{z}(\tau)\}_{\tau \geqslant 0}$ that is initialized at $\mathbf{z}(0) = \mathbf{x}(t)$ and moves inside $S^i(\mathbf{x}(t))$ in the direction captured in the following definition.

Definition 16 Let $i \in [n]$, $S \in \{s_1, \dots, s_m\} \subseteq [i-1]$, and $\mathbf{x} \in [0,1]^n$. Unit vector $\mathbf{d} \in \mathbb{R}^n$ is a direction at point $\mathbf{x}$ if:

- $d_j = 0$, for all $j \notin S \cup \{i\}$, and

- $\langle \nabla V_j(\mathbf{x}), \mathbf{d} \rangle = 0$, for all $j \in S$, and

- the sign of $\begin{vmatrix} \frac{\partial V_{s_1}(\mathbf{x})}{\partial x_{s_1}} & \frac{\partial V_{s_2}(\mathbf{x})}{\partial x_{s_1}} & \cdots & \frac{\partial V_{s_m}(\mathbf{x})}{\partial x_{s_1}} & d_{s_1} \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ \frac{\partial V_{s_1}(\mathbf{x})}{\partial x_{s_m}} & \frac{\partial V_{s_2}(\mathbf{x})}{\partial x_{s_m}} & \cdots & \frac{\partial V_{s_m}(\mathbf{x})}{\partial x_{s_m}} & d_{s_m} \\ \frac{\partial V_{s_1}(\mathbf{x})}{\partial x_i} & \frac{\partial V_{s_2}(\mathbf{x})}{\partial x_i} & \cdots & \frac{\partial V_{s_m}(\mathbf{x})}{\partial x_i} & d_i \end{vmatrix}$ equals the sign of $(-1)^{|S|}$.

If there is a unique unit direction satisfying the above constraints, it is denoted by $D^i_S(\mathbf{x})$. We will also need the notion of exit points.

Definition 17 Suppose $i \in [n]$, $S \subseteq [i-1]$, and $\mathbf{x}' \in [0,1]^n$ is a point where coordinates in S are zero-satisfied and coordinates in $[i-1] \setminus S$ are boundary-satisfied. Then $\mathbf{x}'$ is an *exit point* for epoch (i, S) if it satisfies one of the following:

- Good exit point: Coordinate i is satisfied at $\mathbf{x}'$.
- Bad exit point: For some $j \in S \cup \{i\}$, it holds that $((D^i_S(\mathbf{x}'))_j > 0$ and $x'_j = 1)$, or $((D^i_S(\mathbf{x}'))_j < 0$ and $x'_j = 0)$.
- Middling exit point: For some $j \in [i-1] \setminus S$ it holds that $(V_j(\mathbf{x}') = 0)$ and $((\langle \nabla V_j(\mathbf{x}'), D^i_S(\mathbf{x}') \rangle > 0$ and $x'_j = 0)$ or $(\langle \nabla V_j(\mathbf{x}'), D^i_S(\mathbf{x}') \rangle < 0$ and $x'_j = 1))$.

Algorithm 1 : STON'R (Continuous Dynamics)

Initially $\mathbf{x}(0) \leftarrow (0, \ldots, 0)$, $i \leftarrow 1$, $S \leftarrow \emptyset$, $t \leftarrow 0$.
while $\mathbf{x}(t)$ is not a $\mathrm{VI}(V, [0,1]^n)$ solution **do**
  Initialize epoch (i, S)'s continuous-time dynamics, $\dot{\mathbf{z}}(\tau) = D_s^i(\mathbf{z}(\tau))$, at $\mathbf{z}(0) = \mathbf{x}(t)$.
  while $\mathbf{z}(\tau)$ is not an exit point as in Definition 17 **do**
    Execute $\dot{\mathbf{z}}(\tau) = D_s^i(\mathbf{z}(\tau))$ forward in time.
  end while
  Set $\mathbf{x}(t+\tau) = \mathbf{z}(\tau)$ for all $\tau \in [0, \tau_{exit}]$ (τ_{exit} is the time $\mathbf{z}(t)$ became an exit point).
  if $\mathbf{x}(t+\tau_{exit})$ is a (good exit point) as in Definition 17 **then**
    if i is zero-satisfied at $\mathbf{x}(t+\tau_{exit})$ **then**
      Update $S \leftarrow S \cup \{i\}$.
    end if
    Update $i \leftarrow i + 1$.
  else if $\mathbf{x}(t+\tau_{exit})$ is a (bad exit point) as in Definition 17 for $j = i$ **then**
    Update $i \leftarrow i - 1$ and $S \leftarrow S \setminus \{i-1\}$.
  else if $\mathbf{x}(t+\tau_{exit})$ is a (bad exit point) as in Definition 17 for $j \neq i$ **then**
    Update $S \leftarrow S \setminus \{j\}$.
  else if $\mathbf{x}(t+\tau_{exit})$ is a (middling exit point) as in Definition 17 for $j < i$ **then**
    Update $S \leftarrow S \cup \{j\}$.
  end if
  Set $t \leftarrow t + \tau_{exit}$
end while
return $\mathbf{x}(t)$.

The following state diagram shows the high level overview of the algorithm's behaviour for two variables. It focuses solely on individual epochs and transitions between them, which are determined by the various exit points outlined in Definition 17.

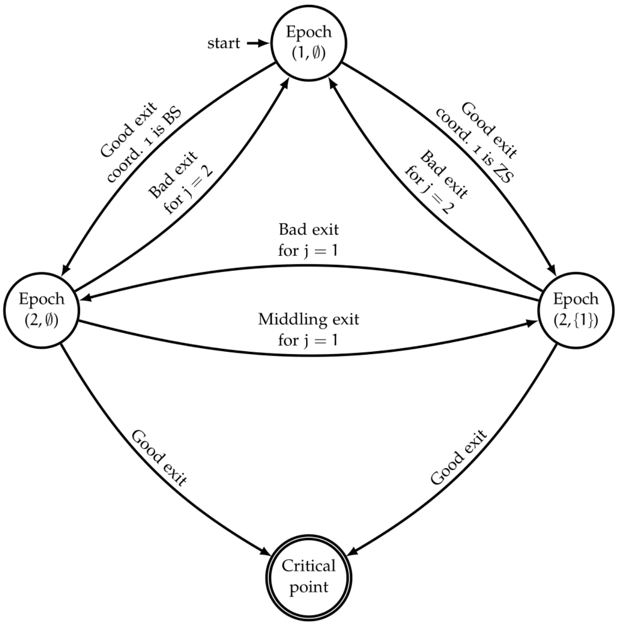

Figure 3.3: Schematic of the STON'R for two variables

3.2 Discrete dynamics

While the continuous dynamics targets exact solutions to $\mathrm{VI}(V, [0, 1]^n)$, the discrete dynamics targets an α-approximate solution to $\mathrm{VI}(V, [0, 1]^n)$ (see Definition 12). In Proposition 9, we stated that the α-approximate solution $\mathbf{x}^*$ to $\mathrm{VI}(V, [0, 1]^n)$ is an $\sqrt{\alpha}$-approximate fixed point of the PGDA dynamics. According to [1, Theorem 5.1], an $\sqrt{\alpha}$-approximate fixed point corresponds to the (ϵ, δ)-local min-max equilibrium (see Definition 7) as follows:

1. Choose $\epsilon > 0$ and $0 < \delta < \sqrt{\frac{2\epsilon}{L}}$.

2. If $\sqrt{\alpha} \leqslant \frac{\sqrt{(G+\delta)^2+4(\epsilon-\frac{L}{2}\delta^2)}-(G+\delta)}{2}$, then $\mathbf{x}^*$ is also an (ϵ, δ)-local min-max equilibrium of f.

When $\delta = \sqrt{\frac{2\epsilon}{L}}$, substituting it into item 2 yields $\sqrt{\alpha} = 0$, indicating that the fixed point should be exact.

Let us now describe the discrete dynamics. In comparison to Algorithm 1, the updates in line 5 do not happen in continuous time but in discrete steps $\mathbf{z}^{(k+1)} \leftarrow \mathbf{z}^{(k)} + D_i^S(\mathbf{z}^{(k)})$. Then there is also a need to revise the definition of the exit points.

Definition 18 Let $\epsilon, \gamma > 0$. Suppose $i \in [n]$, $S \subseteq [i-1]$ and $\mathbf{x}' \in [0,1]^n$ is a point where coordinates in S are zero-satisfied and coordinates in $[i-1] \setminus S$ are boundary-satisfied. Then $\mathbf{x}'$ is an (ϵ, γ)-*exit Point* for epoch (i, S) if it satisfies one of the following:

- Good exit point: Coordinate i is *almost satisfied* at $\mathbf{x}'$, i.e., $|V_i(\mathbf{x}')| \leqslant \epsilon$, or $x'_i = 0$ and $V_i(\mathbf{x}') < \epsilon$, or $x'_i = 1$ and $V_i(\mathbf{x}') > -\epsilon$.

- Bad exit point: For some $j \in S \cup \{i\}$, it holds that $(D_S^i(\mathbf{x}'))_j > 0$ and $x'_j = 1$, or $(D_S^i(\mathbf{x}'))_j < 0$ and $x'_j = 0$.

- Middling exit point: Let $\mathbf{x}'' = \mathbf{x}' + \gamma D_S^i(\mathbf{x}')$ and for some $j \in [i-1] \setminus S$, one of the following holds: $V_j(\mathbf{x}'') > 0$ and $x'_j = 0$, or $V_j(\mathbf{x}'') < 0$ and $x'_j = 1$.

Then the dynamics is updated as follows:

Algorithm 2 : STON'R (Discrete Dynamics)

Initially $\mathbf{x}(0) \leftarrow (0, \ldots, 0)$, $i \leftarrow 1$, $S \leftarrow \emptyset$, $t \leftarrow 0$, $m \leftarrow 0$.
while $\mathbf{x}^{(m)}$ is not an α-approximate VI$(V, [0,1]^n)$ solution **do**
  $\mathbf{z}^{(0)} \leftarrow \mathbf{x}^{(m)}$, $k \leftarrow 0$
  while $\Pi_{[0,1]^n}(\mathbf{z}^{(k)})$ is not an (ϵ, δ)-exit point as in Definition 18 **do**
    $\mathbf{z}^{(k+1)} \leftarrow \mathbf{z}^{(k)} + \gamma \cdot D_S^i(\mathbf{z}^{(k)})$
    $k \leftarrow k+1$
  end while
  $\mathbf{x}^{(m+1)} \leftarrow \Pi_{[0,1]^n}(\mathbf{z}^{(k)})$
  if $\mathbf{x}^{(m+1)}$ is a (good exit point) as in Definition 18 **then**
    if i is zero-satisfied at $\mathbf{x}^{(m+1)}$ **then**
      Update $S \leftarrow S \cup \{i\}$.
    end if
    Update $i \leftarrow i+1$.
  else if $\mathbf{x}^{(m+1)}$ is a (bad exit point) as in Definition 18 for $j = i$ **then**
    Update $i \leftarrow i-1$ and $S \leftarrow S \setminus \{i-1\}$.
  else if $\mathbf{x}^{(m+1)}$ is a (bad exit point) as in Definition 18 for $j \neq i$ **then**
    Update $S \leftarrow S \setminus \{j\}$.
  else if $\mathbf{x}^{(m+1)}$ is a (middling exit point) as in
  Definition 18 for $j < i$ **then**
    Update $S \leftarrow S \cup \{j\}$.
  end if
  Set $m \leftarrow m+1$.
end while
return $\mathbf{x}^{(m)}$

The discrete dynamics produces a point where each coordinate is almost satisfied according to Definition 18. Theorem 39 of [2] establishes that for every $\alpha > 0$ there exist constants $\epsilon, \gamma, \overline{M}, K$ such that Algorithm 2 with step size γ and error ϵ finish after $M \leqslant \overline{M}$ iterations of the while loop at line 2 and it holds that $x^{(M)}$ is an α-approximate solution to $\mathrm{VI}(V, [0,1]^n)$ and for every iteration $m \leqslant M$ of the while loop at line 2, the while loop at line 4 does at most K iterations.

3.3 EXTENSION TO HYPERRECTANGLE

One drawback of STON'R is its restriction to operate over the unit hypercube. However, it is possible to extend the constraint set to a general convex set K by following the strategy outlined in the paper [2, Appendix B]. The strategy requires a maping H from the unit hypercube to K that is bijective and smooth, and its inverse is also smooth.

Let us examine the case when K is a *hyperrectangle*, $K = \bigtimes_{i=1}^{n} [a_i, b_i]$ for some real numbers $a_i < b_i$ with $i = 1, \ldots, n$. Let $f\colon K \to \mathbb{R}$ be a function which ensures that assumptions (2) and (3) for V are satisfied. Let $K' = [0,1]^n$ and $H\colon K' \to K$ be the affine mapping

$$H(\mathbf{x}) = (c_1 x_1 + a_1, \ldots, c_n x_n + a_n),$$

where $\mathbf{x} \in K'$ and $c_i = b_i - a_i$. Define $g\colon K' \to \mathbb{R}$ as $g = f \circ H$. For some $1 \leqslant \ell < n$ define the map $\hat{f} = (\hat{f}_1, \ldots, \hat{f}_n)\colon K \to \mathbb{R}^n$ by

$$\hat{f}_i(\mathbf{x}) = \begin{cases} -\frac{\partial f}{\partial x_i}(\mathbf{x}) & i \leqslant \ell \\ \frac{\partial f}{\partial x_i}(\mathbf{x}) & i > \ell. \end{cases} \qquad i \in [n],\ \mathbf{x} \in K.$$

It follows from the chain rule that, for every $\mathbf{x} \in K'$,

$$\begin{aligned} g'(\mathbf{x}) &= f'(H(\mathbf{x}))H'(\mathbf{x}) = f'(H(\mathbf{x}))\mathrm{diag}(c_1, \ldots, c_n) \\ &= \left(c_1 \frac{\partial f}{\partial x_1}(H(\mathbf{x})), \ldots, c_n \frac{\partial f}{\partial x_n}(H(\mathbf{x})) \right). \end{aligned}$$

Further, define $\hat{g} = (\hat{g}_1, \ldots, \hat{g}_n)$ analogously as

$$\hat{g}_i(\mathbf{x}) = \begin{cases} -c_i \frac{\partial f}{\partial x_i}(H(\mathbf{x})) & i \leqslant \ell \\ c_i \frac{\partial f}{\partial x_i}(H(\mathbf{x})) & i > \ell. \end{cases} \qquad i \in [n],\ \mathbf{x} \in K'.$$

Then, for each $i \in [n]$ and every $\mathbf{x} \in K'$, we get

$$\hat{g}_i(\mathbf{x}) = c_i \hat{f}_i(H(\mathbf{x})). \tag{4}$$

Proposition 13 Let $\mathbf{x} \in K'$. The following are equivalent.

1. $\mathbf{x}$ is a solution to $\mathrm{VI}(\hat{g}, K')$.
2. $H(\mathbf{x})$ is a solution to $\mathrm{VI}(\hat{f}, K)$.
3. $\mathbf{x} = \Pi_{K'}(\mathbf{x} + \hat{g}(\mathbf{x}))$.
4. $H(\mathbf{x}) = \Pi_K(H(\mathbf{x}) + \hat{f}(H(\mathbf{x})))$.
5. $\mathbf{x}$ is a solution to $\mathrm{VI}(G(\mathbf{x}) - \mathbf{x}, K')$, where $G(\mathbf{x}) = H^{-1}(\Pi_K(H(\mathbf{x}) + \hat{f}(H(\mathbf{x}))))$ and $\mathbf{x} \in K'$.

Proof 1. $\Rightarrow$ 2. We need to prove that $\hat{f}(H(\mathbf{x}))^T(H(\mathbf{x}) - \mathbf{y}) \geqslant 0$ for all $\mathbf{y} \in K$. Using (4) the left-hand side equals

$$\sum_{i=1}^{n} \hat{f}_i(H(\mathbf{x}))(H_i(\mathbf{x}) - y_i) = \sum_{i=1}^{n} \frac{\hat{g}_i(\mathbf{x})}{c_i}(H_i(\mathbf{x}) - y_i)$$
$$= \sum_{i=1}^{n} \hat{g}_i(\mathbf{x})(x_i + \frac{a_i - y_i}{c_i}).$$

Let $\mathbf{z} \in \mathbb{R}^n$ be vector with coordinates $z_i = \frac{y_i - a_i}{c_i}$, where $0 \leqslant z_i \leqslant 1$. Then the last term above equals $\hat{g}(\mathbf{x})^T(\mathbf{x} - \mathbf{z})$ and it is nonnegative by the assumption 1. Implication 2. $\Rightarrow$ 1. is proved similarly. The equivalences 1. $\Leftrightarrow$ 3. and 2. $\Leftrightarrow$ 4. follow from Proposition 7. Implication 4. $\Rightarrow$ 5. is trivial since item 4. implies $G(\mathbf{x}) - \mathbf{x} = 0$ for all $\mathbf{x} \in K'$. Let 5. be satisfied. Then pick $\mathbf{y} = G(\mathbf{x})$ in $\mathrm{VI}(G(\mathbf{x}) - \mathbf{x}, K')$, so that $0 \leqslant (G(\mathbf{x}) - \mathbf{x})^T(\mathbf{x} - G(\mathbf{x})) = -\|G(\mathbf{x}) - \mathbf{x}\|^2$. Hence, $G(\mathbf{x}) = \mathbf{x}$, which implies 4.

The first two equivalences in Proposition 13 show how to compute the solution to $\mathrm{VI}(\hat{f}, K)$ by reducing it to computing a solution to $\mathrm{VI}(\hat{g}, K')$. Indeed, any solution $\mathbf{x} \in K'$ to $\mathrm{VI}(\hat{g}, K')$ can be transformed by H to the solution $H(\mathbf{x}) \in K$ for $\mathrm{VI}(\hat{f}, K)$.

3.4 IMPLEMENTATION

We developed the first publicly available implementation of the discrete dynamics (Algorithm 2) of the STON'R algorithm. The algorithm was implemented in Julia (version 1.9) and the source code can be found at

https://gitlab.fel.cvut.cz/kosohmar/StayOnTheRidge.jl.

The main function `run_dynamics(conf::Config)` is defined in the file `src/algorithm.jl`. The function takes a `Config` object as an input, which contains all the necessary information for executing the dynamics, including the function, number of variables, and approximation parameters.

The `Config` type is the supertype of `Config_FD` and `Config_sym`. These two configurations specify how the differentiation for computing the gradient and Hessian matrix will be performed. `Config_FD` enables automatic differentiation using `ForwardDiff.jl`, while `Config_sym` enables symbolic differentiation using `Symbolics.jl`. Symbolic differentiation has the advantage that the gradient and Hessian are computed once before the dynamics execution. The symbolic regime is convenient for debugging the code and observing the behavior of the algorithm with simple functions. However, it is not practical for processing complex functions with a large number of parameters. Hence, automatic differentiation is also available through `Config_FD`. `Config_sym` and `Config_FD` are defined in `src/config.jl` file.

The function `compute_direction(point,i,S,conf::Config)`, which is defined in `src/algorithm.jl`, is important because it is called multiple times. It implements the direction computation according to Definition 16. Certain rows and columns are removed from the transposed Hessian matrix to obtain a new matrix. Next, the nullspace of this matrix is computed using the `LinearAlgebra.jl` package. If the resulting subspace has dimension 1, the direction is uniquely determined, and its polarity is determined by computing the sign of the determinant of this matrix.

The `examples` folder contains numerous example functions, along with their results and sources. In the `test` folder, there are unit tests for the package, covering both the algorithm's functionality and incorrect user inputs. More details can be found on GitLab.

4 NUMERICAL EXPERIMENTS

4.1 MATRIX GAMES

Many algorithms designed to find Nash equilibria are analysed on their performance for matrix games. This is because matrix games are commonly used and easy to create. The loss function for Min-player is a multilinear function defined over the cartesian product of simplices, also known as a simplotope. Since the multilinear function is convex-concave, the Nash equilibrium corresponds to the min-max critical point [11]. Additionally, all Nash equilibria share the same loss value.

Using STON'R to solve matrix games is not straightforward since it would require a mapping H from the unit hypercube to the simplotope, which can be challenging to find due to the strong assumptions outlined in Subsection 3.3. We may ask if STON'R executed over the unit hypercube can give a solution outside the simplotope. We answer in the affirmative in this section.

Consider a matrix game with the matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$, where the first player (rows) has m pure strategies and the second player (columns) has n pure strategies. Matrix $\mathbf{A}$ has entries which are losses of the first player. The sets of mixed strategies are Δ_m and Δ_n, respectively. The expected loss of the first player is determined by the multilinear polynomial $l\colon \Delta_m \times \Delta_n \to \mathbb{R}$, where

$$\ell(\hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\omega}}) = \hat{\boldsymbol{\theta}}^{\mathsf{T}} \mathbf{A} \hat{\boldsymbol{\omega}} = \sum_{i=1}^{m} \sum_{j=1}^{n} a_{ij} \theta_i \omega_j, \quad \hat{\boldsymbol{\theta}} \in \Delta_m, \hat{\boldsymbol{\omega}} \in \Delta_n.$$

Since mixed strategies are elements of the standard simplices, we can express $\theta_m = 1 - (\theta_1 + \cdots + \theta_{m-1})$ and $\omega_n = 1 - (\omega_1 + \cdots + \omega_{n-1})$, and consider the multilinear polynomial f with $m + n - 2$ indeterminates $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_{m-1})$, $\boldsymbol{\omega} = (\omega_1, \ldots, \omega_{n-1})$ given by

$$f(\boldsymbol{\theta}, \boldsymbol{\omega}) = \ell((\boldsymbol{\theta}, 1 - (\theta_1 + \cdots + \theta_{m-1})), (\boldsymbol{\omega}, 1 - (\omega_1 + \cdots + \omega_{n-1}))).$$

We can view f as a real function over $K = [0,1]^{m-1} \times [0,1]^{n-1}$, where $\boldsymbol{\theta} \in [0,1]^{m-1}$ and $\boldsymbol{\omega} \in [0,1]^{n-1}$. Observe that such $\boldsymbol{\theta}$ corresponds to a mixed strategy $\hat{\boldsymbol{\theta}} \in \Delta_m$ iff $\|\boldsymbol{\theta}\|_1 \leqslant 1$. Let $\Theta = \{\boldsymbol{\theta} \in [0,1]^{m-1} \mid \|\boldsymbol{\theta}\|_1 \leqslant 1\}$, $\Omega = \{\boldsymbol{\omega} \in [0,1]^{n-1} \mid \|\boldsymbol{\omega}\|_1 \leqslant 1\}$. Finding a solution over the simplotope is now reduced to finding a solution over the polytope $\Theta \times \Omega$. Executing STON'R over $\Theta \times \Omega$ would require a mapping H from the unit hypercube K to the polytope $\Theta \times \Omega$.

We tried to run STON'R over K. A point in $\Theta \times \Omega$ returned by STON'R is a solution to a given matrix game. This is because it lies in the correct set, and as an equilibrium over K, it is also an equilibrium over $\Theta \times \Omega$. As we mentioned, STON'R can return a point outside $\Theta \times \Omega$. We compare the results given by STON'R and LP for this case. We used LP solver [20].

Example 3 STON'R gives a lower loss than LP.

$$A = \begin{bmatrix} 8.3 & 3.3 \\ 4.77 & 5.92 \\ 8.1 & 4.13 \end{bmatrix}$$

Method	$\boldsymbol{\theta}$	$\boldsymbol{\omega}$	$f(\boldsymbol{\theta}, \boldsymbol{\omega})$
STON'R	$(1, 0.98)$	(0.35)	5.05
LP	$(0.19, 0.81)$	(0.43)	5.43

Table 4.1: Comparison of the STON'R and LP

Example 4 LP gives a lower loss than STON'R.

$$A = \begin{bmatrix} 4.39 & 7.15 & 2.23 \\ 2.94 & 6.72 & 5.9 \\ 4.99 & 6.46 & 3.02 \end{bmatrix}$$

Method	$\boldsymbol{\theta}$	$\boldsymbol{\omega}$	$f(\boldsymbol{\theta}, \boldsymbol{\omega})$
STON'R	$(0, 0.4)$	$(0.05, 1)$	6.56
LP	$(0, 0)$	$(0, 1)$	6.46

Table 4.2: Comparison of the STON'R and LP

These two examples show that the optimal strategies given by STON'R and LP can differ. This is because LP's constraint set is a polytope, while STON'R's constraint set is a unit hypercube.

We compare three methods: STay-ON-the-Ridge (STON'R) [2], Double Oracle (DO) [7], and Regularized Nikaido-Isoda Stochastic Gradient Descent (RNI-SGD) [3]. While STON'R and RNI-SGD converge to an approximate solution to $VI(V, K)$, DO converges to an approximate mixed strategy Nash equilibrium (1) with a finite support. We compare the exploitability of the results according to Definition 2. The exploitability of a mixed strategy profile is computed similarly, but the input function is the expected loss for Min-player. We use the following implementations for computations:

- STON'R: https://gitlab.fel.cvut.cz/kosohmar/StayOnTheRidge.jl,
- RNI-SGD: https://gitlab.mff.cuni.cz/pijalekj/rni-sgd-solver,
- DO: https://github.com/sadda/Double_Oracle.

The results of RNI-SGD were provided by Jan Pijálek, a student at the Faculty of Mathematics and Physics, Charles University.

Nonconvex-nonconcave 2D function [21, Example 3]

$$f(\mathbf{x}) = g(x_1) + Ax_1x_2 - g(x_2), \text{ where}$$

$$g(z) = (z+1)(z-1)(z+3)(z-3), \ A = 11$$

$$\mathbf{x} \in [-4, 4]^2$$

STON'R parameters: $\gamma = 10^{-4}$, $\epsilon = 0.1$.

	STON'R	RNI-SGD
$\mathbf{x}^*$	$(-5.28 \cdot 10^{-4}, 1.24 \cdot 10^{-4})$	$(0, 0)$
$e(\mathbf{x}^*)$	50.014	50

Table 4.3: STON'R and RNI results

	DO
p^*	$0.5 \cdot \delta_{-2.236} + 0.5 \cdot \delta_{2.236}$
q^*	$0.5 \cdot \delta_{-2.236} + 0.5 \cdot \delta_{2.236}$
$e(p^*, q^*)$	$7.01 \cdot 10^{-8}$

Table 4.4: DO results

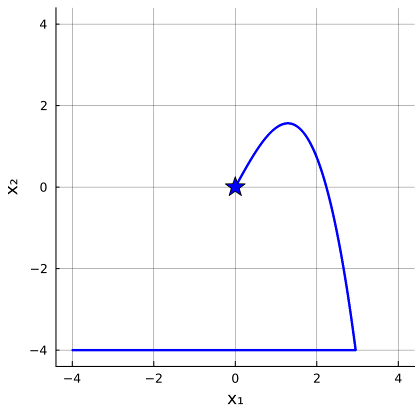

Figure 4.1: STON'R trajectory

The point $(0, 0)$ is the only point with zero gradient. It is not a local min-max equilibrium since the point 0 is a local maximum in x_1 and a local minimum in x_2.

Nonconvex-nonconcave 2D function [21, Example 2]

$$f(\mathbf{x}) = x_1^2 + 3\sin^2(x_1)\sin^2(x_2) - 4x_2^2 - 10\sin^2(x_2)$$

$$\mathbf{x} \in [-10, 10]^2$$

STON'R parameters: $\gamma = 0.001$, $\epsilon = 0.01$.

	STON'R	RNI-SGD
$\mathbf{x}^*$	$(7.11 \cdot 10^{-15}, 7.11 \cdot 10^{-15})$	$(3.05 \cdot 10^{-33}, 0)$
$e(\mathbf{x}^*)$	$-1.15 \cdot 10^{-27}$	$9.29 \cdot 10^{-66}$

Table 4.5: STON'R and RNI results

	DO
p^*	δ_0
q^*	δ_0
$e(p^*, q^*)$	0

Table 4.6: DO results

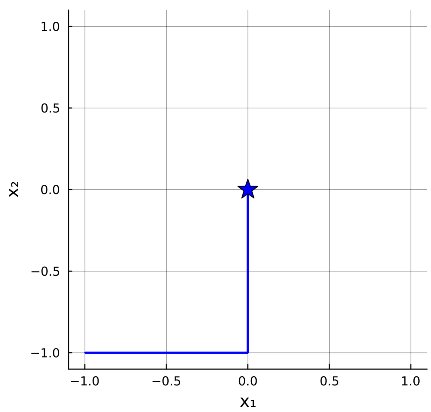

Figure 4.2: STON'R trajectory

The point $(0, 0)$ is the only saddle point of f.

Convex-nonconcave 2D function [22, Example 17]

$$f(\mathbf{x}) = 2x_1^2 + x_2^2 + 4x_1x_2 + \frac{4}{3}x_2^3 - \frac{1}{4}x_2^4$$

$$\mathbf{x} \in [-1, 1]^2$$

STON'R parameters: $\gamma = 0.001$, $\epsilon = 0.01$.

	STON'R	RNI-SGD
$\mathbf{x}^*$	$(0.0016, -0.0016)$	$(0, 0)$
$e(\mathbf{x}^*)$	2.096	2.083

Table 4.7: STON'R and RNI results

	DO
p^*	$0.822 \cdot \delta_{-0.308} + 0.178 \cdot \delta_{-0.315}$
q^*	$0.609 \cdot \delta_1 + 0.391 \cdot \delta_{-0.761}$
$e(p^*, q^*)$	$2.51 \cdot 10^{-5}$

Table 4.8: DO results

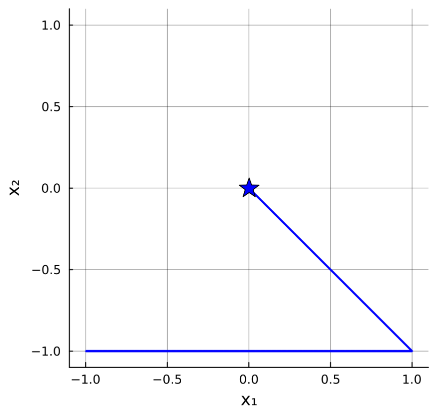

Figure 4.3: STON'R trajectory

The point 0 is a global minimum in x_1 and a local minimum in x_2.

Nonconvex-concave 2D function [5, Example 3]

$$f(\mathbf{x}) = x_1^3 + 2x_1x_2 - x_2^2$$

$$\mathbf{x} \in [-1, 1]^2$$

STON'R parameters: $\gamma = 10^{-4}$, $\epsilon = 10^{-4}$.

	STON'R	RNI-SGD
$\mathbf{x}^*$	$(-1, -1)$	$(-1, -1)$
$e(\mathbf{x}^*)$	2.089	2.089

Table 4.9: STON'R and RNI results

	DO
p^*	$0.583 \cdot \delta_{-1} + 0.417 \cdot \delta_{0.5}$
q^*	$0.512 \cdot \delta_{-0.37} + 0.488 \cdot \delta_{-0.38}$
$e(p^*, q^*)$	$2.19 \cdot 10^{-5}$

Table 4.10: DO results

The point -1 is a local minimum in x_1 (but not global, which increases the exploitability) and a global maximum in x_2.

Nonconvex-nonconcave 6D function [23, Example 6.3 i]

$$f(\mathbf{x}) = \sum_{i=1}^{3}(x_i + x_{3+i}) - \prod_{i=1}^{3}(x_i - x_{3+i})$$

$$\mathbf{x} \in [-1, 1]^6$$

Coordinates 1-3 are minimizing and coordinates 4-6 are maximizing. STON'R parameters: $\gamma = 10^{-4}$, $\epsilon = 0.1$.

	STON'R	RNI-SGD
$\mathbf{x}^*$	$(-1, -1, 1, 1, 1, 1)$	$(1.5 \cdot 10^{-8}, 1.5 \cdot 10^{-8}, 1.5 \cdot 10^{-8}, 1, 1, 1)$
$e(\mathbf{x}^*)$	0	2

Table 4.11: STON'R and RNI results

	DO
p^*	$\delta_{(-1,-1,1)}$
q^*	$\delta_{(1,1,1)}$
$e(p^*, q^*)$	0

Table 4.12: DO results

The point $(-1, -1, 1, 1, 1, 1)$ is a (not unique) saddle point of f, while the point $(0, 0, 0, 1, 1, 1)$ is only min-max critical.

Nonconvex-nonconcave 6D function [23, Example 6.3 ii]

$$f(\mathbf{x}) = -x_1^2 - x_2^2 - x_3^2 + x_4^2 + x_5^2 + x_6^2 + x_1x_5 - x_2x_4 + x_1x_6 - x_3x_4 + x_2x_6 - x_3x_5$$

$$\mathbf{x} \in [-1, 1]^6$$

Coordinates 1-3 are minimizing and coordinates 4-6 are maximizing.
STON'R parameters: $\gamma = 10^{-4}$, $\epsilon = 0.01$.

	STON'R	RNI-SGD
$\mathbf{x}^*$	$(-1,-1,-1,-1,-1,-1)$	$(-1,1,-1,1,-1,1)$
$e(\mathbf{x}^*)$	8	0

Table 4.13: STON'R and RNI results

	DO
p^*	$\delta_{(-0.004,0,0.004)}$
q^*	$\delta_{(0.002,0.005,0.002)}$
$e(p^*, q^*)$	10^{-4}

Table 4.14: DO results

The point $(-1,1,-1,1,-1,1)$ is a (not unique) saddle point of f, while the point $(-1,-1,-1,-1,-1,-1)$ is only min-max critical.

Convex-nonconcave 2D function [24, Example 1]

$$f(\mathbf{x}) = -2x_1x_2^2 + x_1^2 + x_2$$

$$\mathbf{x} \in [-1, 1]^2$$

STON'R parameters: $\gamma = 0.001$, $\epsilon = 0.01$.

	STON'R	RNI-SGD
$\mathbf{x}^*$	$(0.394, 0.632)$	$(0.397, 0.630)$
$e(\mathbf{x}^*)$	$1.05 \cdot 10^{-7}$	≈ 0

Table 4.15: STON'R and RNI results

	DO
p^*	$0.736 \cdot \delta_{0.398} + 0.264 \cdot \delta_{0.395}$
q^*	$0.526 \cdot \delta_{0.631} + 0.474 \cdot \delta_{0.629}$
$e(p^*, q^*)$	$1.83 \cdot 10^{-6}$

Table 4.16: DO results

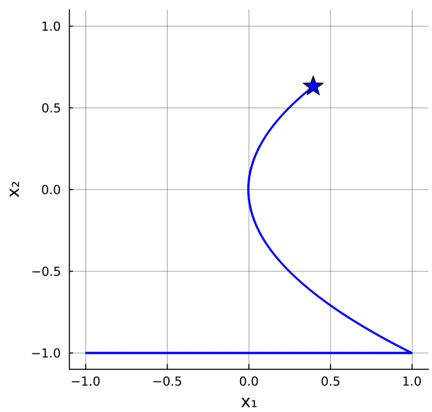

Figure 4.4: STON'R trajectory

The point $(\sqrt[3]{\frac{1}{4}}, \sqrt[3]{\frac{1}{4^2}})$ is the only saddle point of f.

Nonconvex-nonconcave 2D function [25, Figure 1]

$$f(\mathbf{x}) = (x_1 - 0.5)(x_2 - 0.5) + \frac{1}{3}\exp\left(-\left(x_1 - \frac{1}{4}\right)^2 - \left(x_2 - \frac{3}{4}\right)^2\right)$$

$$\mathbf{x} \in [0,1]^2$$

STON'R parameters: $\gamma = 0.001$, $\epsilon = 0.01$.

	STON'R	RNI-SGD
$\mathbf{x}^*$	$(0.419, 0.607)$	$(0.403, 0.597)$
$e(\mathbf{x}^*)$	0.056	0.074

Table 4.17: STON'R and RNI results

	DO
p^*	$0.707 \cdot \delta_1 + 0.293 \cdot \delta_{-1}$
q^*	$0.896 \cdot \delta_{0.445} + 0.104 \cdot \delta_{0.452}$
$e(p^*, q^*)$	$1.65 \cdot 10^{-6}$

Table 4.18: DO results

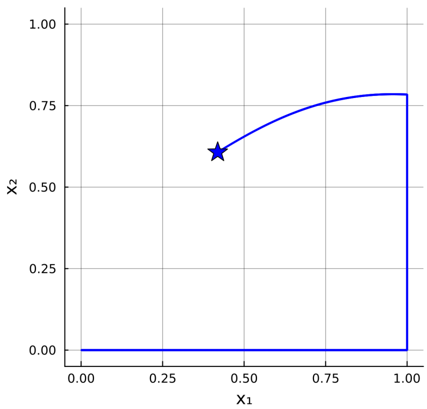

Figure 4.5: STON'R trajectory

The point $\approx (0.4, 0.6)$ is not a saddle point of the function f according to Definition 1. However, it does satisfy 1 when considering a rotated coordinate system.

Convex-nonconcave 2D function [25, Example 2.2]

$$f(\mathbf{x}) = (x_1^4 x_2^2 + x_1^2 + 1)(x_1^2 x_2^4 - x_2^2 + 1)$$

$$\mathbf{x} \in [-1, 1]^2$$

STON'R parameters: $\gamma = 0.001$, $\epsilon = 0.01$.

	STON'R	RNI-SGD
$\mathbf{x}^*$	$(-0.002, -0.002)$	$(0, 0)$
$e(\mathbf{x}^*)$	$8 \cdot 10^{-6}$	0

Table 4.19: STON'R and RNI results

	DO
p^*	δ_0
q^*	δ_0
$e(p^*, q^*)$	0

Table 4.20: DO results

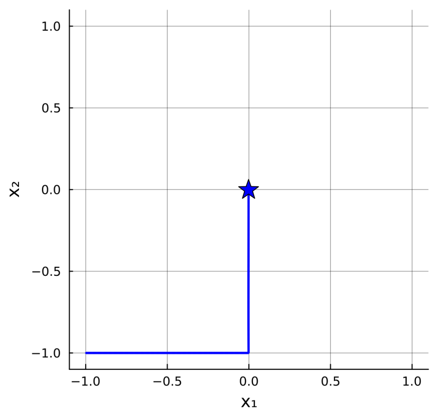

Figure 4.6: STON'R trajectory

The point $(0, 0)$ is the only saddle point of f.

5 CONCLUSION

5.1 STON'R implementation and extensions

We developed an implementation of the STON'R algorithm in Julia. We run STON'R over a wide range of examples, available on GitLab: https://gitlab.fel.cvut.cz/kosohmar/StayOnTheRidge.jl.
Moreover, we formulated an extension of STON'R to the hyperrectangle and discussed the challenge of extending the algorithm to the more general convex sets. We tried to run STON'R over the unit hypercube to solve matrix games and discovered that it may produce results which don't represent probabilities. Building upon this work could involve exploring the potential of applying STON'R to large matrix games, and more broadly, to games with simplex strategy space for each player.

5.2 Comparison of the methods

We presented results given by three methods: STON'R, RNI-SGD, and DO. The method DO achieved low exploitability across all examples as it is designed to minimize exploitability. The results given by STON'R and RNI-SGD were similar in almost all cases, except the 6D examples. When the methods identified a saddle point, the exploitability of the results was zero. Conversely, if the methods found a min-max critical point that is not a saddle point, the exploitability was nonzero.

5.3 Generalized min-max critical points

We extended the concept of the min-max critical point to locally Lipschitz functions, proved its existence over convex compact set, and showed similar properties that hold for the smooth min-max critical point. This extension is motivated by the practical problems, such as those encountered in deep neural networks, where the loss functions may not be differentiable at a set of points with measure zero. Further research could explore the game-theoretical perspective of the generalized min-max critical points, investigating whether they represent an approximate local min-max equilibrium, or developing methods for finding these points.

BIBLIOGRAPHY

[1] 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*. 2021, pp. 1466–1478.

[2] Constantinos Daskalakis, Noah Golowich, Stratis Skoulakis, and Emmanouil Zampetakis. "Stay-on-the-ridge: Guaranteed convergence to local minimax equilibrium in nonconvex-nonconcave games." In: *The Thirty Sixth Annual Conference on Learning Theory*. PMLR. 2023, pp. 5146–5198.

[3] Ioannis Tsaknakis and Mingyi Hong. "Finding first-order nash equilibria of zero-sum games with the regularized nikaido-isoda function." In: *International Conference on Artificial Intelligence and Statistics*. PMLR. 2021, pp. 1189–1197.

[4] Maher Nouiehed, Maziar Sanjabi, Tianjian Huang, Jason D Lee, and Meisam Razaviyayn. "Solving a class of non-convex minmax games using iterative first order methods." In: *Advances in Neural Information Processing Systems* 32 (2019).

[5] Meisam Razaviyayn, Tianjian Huang, Songtao Lu, Maher Nouiehed, Maziar Sanjabi, and Mingyi Hong. "Nonconvex min-max optimization: Applications, challenges, and recent theoretical advances." In: *IEEE Signal Processing Magazine* 37.5 (2020), pp. 55–66.

[6] Constantinos Daskalakis, Paul W Goldberg, and Christos H Papadimitriou. "The complexity of computing a Nash equilibrium." In: *Communications of the ACM* 52.2 (2009), pp. 89–97.

[7] Lukáš Adam, Rostislav Horčík, Tomáš Kasl, and Tomáš Kroupa. "Double oracle algorithm for computing equilibria in continuous games." In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 35. 6. 2021, pp. 5070–5077.

[8] J v. Neumann. "Zur theorie der gesellschaftsspiele." In: *Mathematische annalen* 100.1 (1928), pp. 295–320.

[9] Irving L Glicksberg. "A further generalization of the Kakutani fixed theorem, with application to Nash equilibrium points." In: *Proceedings of the American Mathematical Society* 3.1 (1952), pp. 170–174.

[10] Chi Jin, Praneeth Netrapalli, and Michael Jordan. "What is local optimality in nonconvex-nonconcave minimax optimization?" In: *International conference on machine learning*. PMLR. 2020, pp. 4880–4889.

[11] R Tyrrell Rockafellar and Roger J-B Wets. *Variational analysis*. Vol. 317. Springer Science & Business Media, 2009.

[12] Francisco Facchinei and Jong-Shi Pang. *Finite-dimensional variational inequalities and complementarity problems*. Springer, 2003.

[13] Giancarlo Bigi, Lorenzo Lampariello, Simone Sagratella, and Valerio Giuseppe Sasso. "Approximate variational inequalities and equilibria." In: *Computational Management Science* 20.1 (2023), p. 43.

[14] Francis Clarke. *Functional analysis, calculus of variations and optimal control*. Vol. 264. Springer, 2013.

[15] SC Fang and EL Peterson. "Generalized variational inequalities." In: *Journal of Optimization Theory and Applications* 38 (1982), pp. 363–383.

[16] Frank H Clarke. "Generalized gradients and applications." In: *Transactions of the American Mathematical Society* 205 (1975), pp. 247–262.

[17] A Hitchhiker's Guide. *Infinite dimensional analysis*. Springer, 2006.

[18] Frank H Clarke. *Optimization and nonsmooth analysis*. SIAM, 1990.

[19] Tomáš Kroupa and Ondrej Majer. "Optimal strategic reasoning with McNaughton functions." In: *International journal of approximate reasoning* 55.6 (2014), pp. 1458–1468.

[20] David Avis, Gabriel D Rosenberg, Rahul Savani, and Bernhard Von Stengel. "Enumeration of Nash equilibria for two-player games." In: *Economic theory* 42 (2010), pp. 9–37. URL: http://banach.lse.ac.uk.

[21] Taoli Zheng, Linglingzhi Zhu, Anthony Man-Cho So, José Blanchet, and Jiajin Li. "Universal Gradient Descent Ascent Method for Nonconvex-Nonconcave Minimax Optimization." In: *Advances in Neural Information Processing Systems* 36 (2024).

[22] Leonard Adolphs, Hadi Daneshmand, Aurelien Lucchi, and Thomas Hofmann. "Local saddle point optimization: A curvature exploitation approach." In: *The 22nd International Conference on Artificial Intelligence and Statistics*. PMLR. 2019, pp. 486–495.

[23] Jiawang Nie, Zi Yang, and Guangming Zhou. "The saddle point problem of polynomials." In: *Foundations of Computational Mathematics* (2021), pp. 1–37.

[24] Tomáš Kroupa and Tomáš Votroubek. "Multiple Oracle Algorithm to Solve Continuous Games." In: *International Conference on Decision and Game Theory for Security*. Springer. 2022, pp. 149–167.

[25] Panayotis Mertikopoulos, Bruno Lecouat, Houssam Zenati, Chuan-Sheng Foo, Vijay Chandrasekhar, and Georgios Piliouras. "Optimistic mirror descent in saddle-point problems: Going the extra (gradient) mile." In: *arXiv preprint arXiv:1807.02629* (2018).

USED SOFTWARE

The following software was used in the development of this thesis:

- ChatGPT (OpenAI)[1] for text feedback and rephrasing suggestions
- Grammarly[2] for grammar and spelling checking

1 https://openai.com/
2 https://www.grammarly.com/