

Learning Variational Inequalities from Data: Fast Generalization Rates under Strong Monotonicity

Eric Zhao*

Tatjana Chavdarova[†]Michael Jordan[‡]

Abstract

Variational inequalities (VIs) are a broad class of optimization problems encompassing machine learning problems ranging from standard convex minimization to more complex scenarios like min-max optimization and computing the equilibria of multi-player games. In convex optimization, *strong convexity* allows for fast statistical learning rates requiring only $\Theta(1/\varepsilon)$ stochastic first-order oracle calls to find an ε -optimal solution, rather than the standard $\Theta(1/\varepsilon^2)$ calls. In this paper, we explain how one can similarly obtain fast $\Theta(1/\varepsilon)$ rates for learning VIs that satisfy *strong monotonicity*, a generalization of strong convexity. Specifically, we demonstrate that standard stability-based generalization arguments for convex minimization extend directly to VIs when the domain admits a small covering, or when the operator is integrable and sub-optimality is measured by potential functions; such as when finding equilibria in multi-player games.

1 Introduction

Variational inequalities (Stampacchia, 1964) model a broad range of optimization and fixed-point problems. They are concerned with finding a solution \mathbf{z}^* from a continuous domain \mathcal{Z} , such that

$$\langle \mathbf{z} - \mathbf{z}^*, F(\mathbf{z}^*) \rangle \geq 0, \quad \forall \mathbf{z} \in \mathcal{Z}, \quad (\text{VI})$$

where $F: \mathcal{Z} \mapsto \mathbb{R}^n$ is a continuous map, and \mathcal{Z} is a subset of the Euclidean space \mathbb{R}^d . The operator F can be understood as encoding both the optimality condition and first-order information of a VI.

Variational inequalities (VIs) are a popular framework primarily due to their generality—they encompass convex minimization problems (when $F \equiv \nabla f$), min-max optimization problems, complementarity problems (Cottle and Dantzig, 1968), and equilibrium computation problems for general games. VIs therefore provide a language for proving results that are applicable to a wide range of optimization problems. As an example, consider a game with k agents, each with a strategy $\mathbf{z}_i \in \mathbb{R}^{d_i}$ and aiming to optimize their own objective $f_i: \mathbb{R}^{d_i} \rightarrow \mathbb{R}$. Denoting the players’ joint strategy with $\mathbf{z} \equiv [\mathbf{z}_1^\top, \dots, \mathbf{z}_k^\top]^\top \in \mathbb{R}^d$, $d = \sum_{i=1}^k d_i$, we can write the problem of finding an equilibrium in this game as equivalently solving a VI whose operator F represents the gradients of each player’s objective:

$$F(\mathbf{z}) \equiv \begin{bmatrix} \nabla_{\mathbf{z}_1} f_1(\mathbf{z}) \\ \dots \\ \nabla_{\mathbf{z}_k} f_k(\mathbf{z}) \end{bmatrix}.$$

*University of California, Berkeley. Email: eric.zh@berkeley.edu

[†]University of California, Berkeley. Email: tatjana.chavdarova@berkeley.edu

[‡]University of California, Berkeley. Email: jordan@cs.berkeley.edu

We refer the reader to Facchinei and Pang (2003) for more examples. Due to their generality, VIs also find applications in economics, game theory, and machine learning, to multi-agent reinforcement learning, GANs, and other problems involving tradeoffs (Nagurney and Zhang, 1995; Martínez-Legaz, 1990; Scutari et al., 2010; Lowe et al., 2017; Sidahmed and Chavdarova, 2024; Gidel et al., 2019).

In many applications of variational inequalities, the operator F may not be known a priori. Instead, one may only have access to historical datapoints for the operator or other noisy finite estimates of operator evaluations. Indeed, this is often the case in many practical applications of VIs: (i) *auction games*—where an auctioneer and bidders interact so as to allocate resources (McAfee and McMillan, 1987; Riley, 1989; Milgrom, 2004; Krishna, 2009); (ii) various *network games* (Galeotti et al., 2010; Parise and Ozdaglar, 2019)—which model the interaction between parts of information systems controlled by different parties; or (iii) *alternating multi-agent games*—where the goal is to perform well on an unseen game (e.g., mini-games in Vinyals et al., 2017). The key challenge with learning VI solutions in these applications is contending with uncertainty due to noisy observation model. This noise can originate from different sources depending on the application, such as stochasticity in agents’ decisions, noisy mini-batch evaluation of the operator, or having access only to unbiased variations of the “ground-truth” operator F . One approach to learning VI solutions from noisy data is to run a conventional algorithm, such as extragradient descent, on an empirical estimate of the operator and then bound generalization error, i.e., prove that a near-optimal solution for the empirical operator estimate is also near-optimal for the true operator.

One classical approach to bounding generalization error exploits the fact that learning algorithms that exhibit *algorithmic stability* (Rogers and Wagner, 1978; Devroye and Wagner, 1979)—a measure of how sensitive an algorithm’s outputs are to small perturbations in the input dataset—must have limited generalization error. In strongly convex optimization problems, the key advantage of stability-based generalization analysis is that—even for simple first-order methods—they provide tighter sample complexity rates of order $O(1/\varepsilon)$ where ε denotes the error tolerance (Bousquet and Elisseeff, 2002). This is in contrast to the $O(1/\varepsilon^2)$ samples needed by stochastic approximation methods (Nemirovski and Yudin, 1983). Recent papers have extended these stability-based generalization results to strongly-convex strongly-concave zero-sum games (Li and Liu, 2022; Farnia and Ozdaglar, 2021; Zhang et al., 2020).

In this paper, we demonstrate how to directly extend stability-based generalization proofs for convex minimization to variational inequalities. In particular, we demonstrate how fast sample complexity rates of $O(1/\varepsilon)$ are readily obtainable in strongly monotone VIs, which are the natural VI generalization of strongly convex optimization (see Section 2) and include strongly-convex-concave min-max optimization as a special case. A technical issue that we address is that the conventional measure of suboptimality in VIs, the *gap function*, can be unstable. This is a property of gap functions that is not unique to VIs and can be bypassed in two ways. First, if a VI’s domain admits a small covering, the generalization error of the gap function is stable and fast generalization rates can be shown accordingly. Second, if the operator is integrable and suboptimality is measured in terms of potential functions, one can still prove fast generalization rates by analyzing the stability of a surrogate for the gap function; a common example is computing the Nash equilibrium of a multi-player game whose players have strongly convex/concave utilities.

Related works. Algorithmic stability is a standard approach to characterizing the generalization of a learning algorithm (Mukherjee et al., 2006; Shalev-Shwartz et al., 2010; Charles and Papailiopoulos, 2018; Lei and Ying, 2020; Kuzborskij and Lampert, 2017). The work of Bousquet and Elisseeff (2002) derived high-probability generalization bounds, which was further improved on by Feldman

and Vondrak (2019); Klochkov and Zhivotovskiy (2021). These results prove fast sample complexity rates of $O(1/\varepsilon)$ for strongly convex minimization problems. Recently, a number of works have proven similar fast rates for strongly-convex-concave games (Li and Liu, 2022; Farnia and Ozdaglar, 2021; Zhang et al., 2020). This note generalizes and streamlines these results by showing that, in fact, the same stability arguments used in convex minimization can be directly extended to VIs.

We also note that there is a long history of results studying the convergence of first-order methods in VIs. In deterministic problems, the development of efficient algorithms with convergence guarantees has recently been the focus of interest in machine learning and optimization in both unconstrained (see, e.g., Tseng, 1995; Daskalakis et al., 2018; Mokhtari et al., 2019, 2020; Golowich et al., 2020; Azizian et al., 2020; Chavdarova et al., 2023; Gorbunov et al., 2022b; Bot et al., 2022), and constrained settings (Nemirovski and Yudin, 1983; Beck and Teboulle, 2003; Nemirovski, 2004; Cai et al., 2022; Yang et al., 2023). For stochastic problems, the seminal work of Juditsky et al. (2011) on stochastic mirror prox provides results on the stochastic approximation approach; these rates are slower in strongly monotone settings relative to ours, and cover stochastic approximation algorithms whereas our results cover *batch* algorithms.

2 Preliminaries

We will use bold small letters denote vectors and curly capital letters denote sets. We also use \mathcal{Z} to denote a convex and compact set in the Euclidean space \mathcal{E} , with inner product $\langle \cdot, \cdot \rangle$ and norm $\|\cdot\|$.

2.1 Variational inequalities (VIs)

A VI is represented by the tuple (F, \mathcal{Z}) containing its operator F and domain \mathcal{Z} . We use D to denote the diameter of the domain \mathcal{Z} where $D := \max_{\mathbf{z}, \mathbf{z}' \in \mathcal{Z}} \|\mathbf{z} - \mathbf{z}'\|$ will be treated as a constant; for example, if \mathcal{Z} is the probability simplex Δ_d and $\|\cdot\|$ is the ℓ_1 norm, then $D = 2$.

The most common solution concept for a variational inequality is the *Stampacchia solution*, which is defined as a solution $\mathbf{z}_* \in \mathcal{Z}$ where $\langle F(\mathbf{z}_*), \mathbf{z}_* - \mathbf{z} \rangle \leq 0$ for all $\mathbf{z} \in \mathcal{Z}$. A standard quantity for measuring how close a solution $\mathbf{z} \in \mathcal{Z}$ is to being a Stampacchia solution is the *gap function*:

$$\text{Err}_F(\mathbf{z}) \triangleq \max_{\mathbf{u} \in \mathcal{Z}} \langle F(\mathbf{z}), \mathbf{z} - \mathbf{u} \rangle .$$

When F is clear from context, we will omit the subscript and write Err .

When we have a VI with a conservative operator F — F is the gradient(s) of a potential function(s)—another common numerical quantification of suboptimality is the gap between the best attainable potential versus the iterate’s potential, which we will refer to as the *potential gap*. Specifically, if the domain has a product structure $\mathcal{Z} = \mathcal{Z}_1 \times \cdots \times \mathcal{Z}_k$, we can interpret the VI as a game with k players. For each player $i \in [d]$, the restriction of the operator F to subspace \mathcal{Z}_i , denoted by F_i , can be interpreted as the gradient of their potential/utility function f_i . We can then define the suboptimality of a solution $\mathbf{z} \in \mathcal{Z}$ as sum of each player’s suboptimality:

$$\text{Err}_F(\mathbf{z}) \triangleq \sum_{i \in [k]} f_i(\mathbf{z}) - \min_{\mathbf{z}_i^* \in \mathcal{Z}_i} f_i(\mathbf{z}_i^*, \mathbf{z}_{-i}).$$

Throughout, we will adopt the standard assumptions that an operator F is *L-Lipschitz* and *monotone*. These conditions ensure that a (weak) solution to the VI exists (Facchinei and Pang, 2003, Chapter 2).

Assumption 2.1 (*L*-Lipschitz operator). *An operator $F: \mathcal{Z} \rightarrow \mathcal{E}$ is L -Lipschitz if and only if:*

$$\|F(\mathbf{z}) - F(\mathbf{z}')\| \leq L \|\mathbf{z} - \mathbf{z}'\|, \quad \forall \mathbf{z}, \mathbf{z}' \in \mathcal{Z}.$$

Assumption 2.2 (Monotone operator). *An operator $F: \mathcal{Z} \rightarrow \mathcal{E}$ is monotone if and only if:*

$$\langle \mathbf{z} - \mathbf{z}', F(\mathbf{z}) - F(\mathbf{z}') \rangle \geq 0, \quad \forall \mathbf{z}, \mathbf{z}' \in \mathcal{Z}.$$

Monotonicity can be understood as an analog of the assumption of convexity. In some cases, we may assume the operator is strongly monotone, which is a stronger assumption as it implies monotonicity. One can understand strong monotonicity as generalizing strong convexity.

Assumption 2.3 (μ -strongly monotone operator). *An operator $F: \mathcal{Z} \rightarrow \mathcal{E}$ is said to be μ -strongly monotone if and only if there exists $\mu > 0$ such that:*

$$\langle \mathbf{z} - \mathbf{z}', F(\mathbf{z}) - F(\mathbf{z}') \rangle \geq \mu \|\mathbf{z} - \mathbf{z}'\|^2, \quad \forall \mathbf{z}, \mathbf{z}' \in \mathcal{Z}.$$

In strongly monotone settings, the ratio L/μ is known as the *condition number* of the problem. For ease of presentation, we will focus on how generalization rates scale with dimension, sample complexity, and error tolerance. As such, the condition number should be understood to be a constant quantity on which we do not necessarily seek a tight characterization, although we will always state dependence on the condition number for completeness.

Finally, we will assume that the VIs we study are in an unconstrained setting.

Assumption 2.4 (Unconstrained setting). *There exists $\mathbf{z}^* \in \mathcal{Z}$ such that $F(\mathbf{z}^*) = 0$.*

2.2 Generalization and algorithmic stability

We are interested in solving a variational inequality problem without direct access to an operator oracle $\mathbf{z} \mapsto F(\mathbf{z})$ by learning from data. In this setting, we are restricted to only knowing the domain \mathcal{Z} and having access to a noisy operator oracle Ξ and a series of noise inputs ζ_1, \dots, ζ_n drawn i.i.d. from some data distribution \mathcal{D} whose support we will denote as \mathcal{X} . Formally, Ξ is a Borel function that given a sample $\zeta \sim \mathcal{D}$ provides the unbiased estimate

$$\forall \mathbf{z} \in \mathcal{Z} : \mathbb{E}_{\zeta \sim \mathcal{D}} [\Xi(\mathbf{z}, \zeta)] = F(\mathbf{z}).$$

Each ζ_i can be understood as a datapoint in a dataset X of size n . Given a dataset $X \in \mathcal{X}^n$, we will use the shorthand $\widehat{F}_X(\mathbf{z}) = \frac{1}{n} \sum_{i=1}^n \Xi(\mathbf{z}, \zeta_i)$ and $\widehat{F}_{X_i}(\mathbf{z}) = \Xi(\mathbf{z}, \zeta_i)$. Two datasets $X, X' \in \mathcal{X}^n$ are *neighboring* if they disagree in at most one entry.

We assume that noisy operator estimates satisfy the same assumptions as the true operator F (i.e., Assumptions 2.2, 2.1), as is standard in generalization literature, e.g. Hardt et al. (2016); Farnia and Ozdaglar (2021). We will also assume that operator samples are of bounded norm to control the variance of the noise.

Assumption 2.5 (Bounded Operator Norm). *The operator F has bounded norm at every $\mathbf{z} \in \mathcal{Z}$:*

$$\|F(\mathbf{z})\|^2 \leq K, \quad \forall \mathbf{z} \in \mathcal{Z}.$$

An empirical variant of Algorithms 1 and 2 then follows from running the algorithms on the empirical operator \widehat{F}_X computed from a dataset $X \in \mathcal{X}^n$ rather than the true operator $F(\mathbf{z})$. We aim to understand the associated generalization error of using our empirical operator \widehat{F}_X instead of

the true one, F . Given a candidate solution $\mathbf{z} \in \mathcal{Z}$ produced by a learning algorithm trained on a dataset \mathcal{X}_n , we define its empirical gap function value as

$$\widehat{\text{Err}}_X(\mathbf{z}) := \max_{\mathbf{w} \in \mathcal{Z}} \langle \widehat{F}_X(\mathbf{z}), \mathbf{z} - \mathbf{w} \rangle$$

and its generalization gap as $\text{Err}(\mathbf{z}) - \widehat{\text{Err}}_X(\mathbf{z})$. In this work, we will mainly focus on settings where we drive the empirical gap function $\widehat{\text{Err}}_X(\mathbf{z})$ to zero by training on the empirical operator \widehat{F}_X for many iterations, and only seek to control the generalization gap. We analogously define the empirical potential gap $\widehat{\text{Err}}_X(\mathbf{z})$ and the generalization potential gap $\text{Err}(\mathbf{z}) - \widehat{\text{Err}}_X(\mathbf{z})$.

Stability. We say that a mapping is *uniformly stable* if changing one record ζ_i in its input (a dataset) does not significantly change the output.

Definition 2.6 (Uniformly Stable Algorithm). *An algorithm $A: \mathcal{X}^n \rightarrow \mathcal{Z}$ is γ -uniformly stable if, for all datasets $X \in \mathcal{X}^n$ and $X' \in \mathcal{X}^n$ that differ in only one element it holds that:*

$$\|A(X) - A(X')\| \leq \gamma.$$

Composing a uniformly stable algorithm with a Lipschitz loss results in a uniformly bounded function. Uniformly bounded functions admit powerful generalization bounds (Bousquet and Elisseeff, 2002) that, for instance, can allow for obtaining dimension-free rates or fast rates.

Definition 2.7 (Uniformly Bounded Functions). *A mapping $f: \mathcal{X}^n \times \mathcal{X} \rightarrow [0, 1]$ is γ -uniformly bounded if, for all datasets $X \in \mathcal{X}^n$ and $X' \in \mathcal{X}^n$ that differ in only one element, and all evaluation data points $\mathbf{x} \in \mathcal{X}$, it holds that:*

$$|f(X, \mathbf{x}) - f(X', \mathbf{x})| \leq \gamma.$$

2.3 VI methods

Among the most commonly used first-order algorithms for solving VIs are the (stochastic) *gradient descent* and the *extragradient* methods (Korpelevich, 1976). Gradient descent naturally extends to the setting of VI optimization as follows:

$$\mathbf{z}_{t+1} = \mathbf{z}_t - \eta F(\mathbf{z}_t), \tag{GD}$$

where $\eta \in [0, 1]$ denotes a step size.

The *extragradient* (EG) method proposed by Korpelevich (1976) uses an extrapolated point $\mathbf{z}_{t+\frac{1}{2}}$ using GD: $\mathbf{z}_{t+\frac{1}{2}} = \mathbf{z}_t - \eta F(\mathbf{z}_t)$, and then adds the gradient at the extrapolated point to the current iterate \mathbf{z}_t :

$$\mathbf{z}_{t+1} = \mathbf{z}_t - \eta F(\mathbf{z}_{t+\frac{1}{2}}) = \mathbf{z}_t - \eta F(\mathbf{z}_t - \eta F(\mathbf{z}_t)). \tag{EG}$$

More detailed descriptions of these methods can be found in Appendix A.1. Both (GD) and (EG) converge for strongly monotone VIs in deterministic (batch) (Korpelevich, 1976; Tseng, 1995) and stochastic settings (Hsieh et al., 2020; Loizou et al., 2021; Gorbunov et al., 2022a). (GD) is often used in practice due to its simplicity, and we study the popular (EG) as handling the extrapolation for sample complexity requires additional instructive steps. The results we demonstrate for (GD) and (EG) can also be applied to obtain corresponding sample complexity bounds for other first-order VI methods, such as *optimistic* GD (Popov, 1980).

Stability. It is known that the iterates produced by gradient descent are stable in norm in settings with smoothness (Assumption 2.1), strong monotonicity (Assumption 2.3) and bounded operator norms (Assumption 2.5) (Hardt et al., 2016).

Lemma 2.8 (GD stability). *Consider a smooth strongly monotone variational inequality (F, \mathcal{Z}) (Assumption 2.1, 2.3, 2.5). Given neighboring datasets $X \in \mathcal{X}_n$ and $X' \in \mathcal{X}_n$, let \mathbf{z}_T and \mathbf{z}'_T denote the T th iterates of gradient descent (Algorithm 1) on the empirical operators given by X and X' respectively. For step sizes $0 < \eta < \frac{2\mu}{L^2}$, uniform stability holds with*

$$\|\mathbf{z}_T - \mathbf{z}'_T\| \leq \frac{2K}{n(2\mu - \eta L^2)}.$$

We can extend this result to argue the stability of extragradient descent. A key step stability results of this kind is the argument that each step of the algorithm is contractive with respect to a fixed operator, which we establish for the extragradient algorithm.

Lemma 2.9 (EG contractiveness). *Let $F: \mathcal{Z} \rightarrow \mathbb{R}^d$ be a L -Lipschitz μ -strongly monotone operator (Assumptions 2.1, 2.3). Let $G_{EG}(\mathbf{z}, F)$ denote the EGD update of an iterate \mathbf{z} given operator F , that is $G_{EG} \equiv I - \eta F \circ (I - \eta F)$. For any learning rate $\eta > 0$, and $\forall \mathbf{z}, \mathbf{z}' \in \mathcal{Z}$ we have:*

$$G_{EG}(\mathbf{z}, F) \leq \underbrace{(1 - \eta + \eta^2 L^2 (1 - 2\eta\mu + \eta^2 L^2))}_{\triangleq c_{EG}(\eta)} \|\mathbf{z} - \mathbf{z}'\|_2^2.$$

That is, G_{EG} is contractive for any step size η where $c_{EG}(\eta) < 1$ and $\{\eta_i \mid c_{EG}(\eta_i) < 1\}$ is non empty.

The stability of extragradient then follows via a standard growth lemma argument.

Lemma 2.10 (EG stability). *Consider a smooth strongly monotone variational inequality (F, \mathcal{Z}) (Assumption 2.1, 2.3, 2.5). Given neighboring datasets $X \in \mathcal{X}_n$ and $X' \in \mathcal{X}_n$, let \mathbf{z}_T and \mathbf{z}'_T denote the T th iterates of extragradient descent (Algorithm 2) on the empirical operators given by X and X' respectively. For step sizes $\eta > 0$, uniform stability holds with*

$$\|\mathbf{z}_T - \mathbf{z}'_T\| \leq \frac{2K}{n(1 + \eta - \eta L^2(1 - \eta)(1 - 2\eta\mu + \eta^2 L^2))(1 - \eta L - (1 + \eta L))}.$$

3 Fast Gap Function Bounds for Small Domains

The primary challenge with performing stability-based generalization analysis for variational inequalities is that the variational term for gap functions, $\arg \max_{\mathbf{z}^* \in \mathcal{Z}} \langle F(\mathbf{z}), \mathbf{z} - \mathbf{z}^* \rangle$, is unstable. This challenge is not unique to variational inequalities but rather an artifact of the gap function. From another perspective, stability proofs involve a symmetrization step that typically assume that the training loss is a linear combination of losses computed on individual datapoints, whereas the term $\arg \max_{\mathbf{z}^* \in \mathcal{Z}} \langle \widehat{F}_X(\mathbf{z}), \mathbf{z} - \mathbf{z}^* \rangle$ may depend nonlinearly on the dataset X .

A natural approach is to take a union bound over the possible values of this variational term. Specifically, we can define for every point in the solution set $w \in \mathcal{Z}$, the loss

$$\langle \widehat{F}(\mathbf{z}), \mathbf{z} - w \rangle = \frac{1}{n} \sum_{i=1}^n \langle \Xi(\mathbf{z}, \zeta_i), \mathbf{z} - w \rangle, \quad (1)$$

which is linear in the datapoints $\{\zeta_i\}_{i \in [n]}$. This reduces the problem of bounding the generalization error of the gap function to bounding the generalization error of a linear loss function (1) and then taking a union bound over the possible values of w in the domain \mathcal{Z} . Indeed, for strongly monotone VIs with small domains, this approach provides fast in-expectation generalization error bounds of $O(1/n)$, where n is the size of the dataset.

Theorem 3.1. *Let (F, \mathcal{Z}) be a smooth strongly monotone VI (Assumptions 2.1, 2.3, 2.5). The expected generalization gap for a γ -stable algorithm $A: \mathcal{X}^n \rightarrow \mathcal{Z}$ given n datapoints is, for $r > 0$,*

$$\mathbb{E}_{X \sim \mathcal{D}^n} [\text{Err}_F(A(X)) - \text{Err}_{\widehat{F}}(A(X))] \leq O\left(Kr + (LD + K)\gamma \log(N(\mathcal{Z}, r, \|\cdot\|))\right),$$

where $N(\mathcal{Z}, r, \|\cdot\|)$ is the r -covering number for \mathcal{Z} .

Proof. Let us first fix a choice of $\mathbf{w} \in \mathcal{Z}$. This stability proof begins with a standard symmetrization argument. Let us define $X^{(i)}$ as a copy of the dataset X but whose i th record X_i is replaced with that of dataset X' , i.e. X'_i . We can thus write

$$\mathbb{E}_{X \sim \mathcal{D}^n} [\langle F(A(X)), A(X) - \mathbf{w} \rangle] = \frac{1}{n} \sum_{\substack{X \sim \mathcal{D}^n \\ X' \sim \mathcal{D}^n}} \mathbb{E} \left[\langle \widehat{F}_{X_i}(A(X^{(i)})), A(X^{(i)}) - \mathbf{w} \rangle \right]. \quad (2)$$

Next, recall that, due to the γ -stability of the algorithm A , for any realization of the datasets X and X' , we have by definition that $\|A(X) - A(X^{(i)})\| \leq \gamma$. Applying Hölder's inequality and the L -smoothness of \widehat{F}_{X_i} , we can bound the symmetrization gap by

$$\begin{aligned} & \left\langle \widehat{F}_{X_i}(A(X^{(i)})), A(X^{(i)}) - \mathbf{w} \right\rangle - \left\langle \widehat{F}_{X_i}(A(X)), A(X) - \mathbf{w} \right\rangle \\ & \leq \left\langle \widehat{F}_{X_i}(A(X^{(i)})) - \widehat{F}_{X_i}(A(X)), A(X^{(i)}) - \mathbf{w} \right\rangle + \left\langle \widehat{F}_{X_i}(A(X)), A(X) - A(X^{(i)}) \right\rangle \\ & \leq LD\gamma + K\gamma. \end{aligned}$$

As a result, the random variable

$$\Delta(X, X', \mathbf{w}) := \frac{1}{n} \sum_{i=1}^n \left\langle \widehat{F}_{X_i}(A(X)), A(X) - \mathbf{w} \right\rangle - \left\langle \widehat{F}_{X_i}(A(X^{(i)})), A(X^{(i)}) - \mathbf{w} \right\rangle,$$

is always bounded by $\Delta(X, X', \mathbf{w}) \leq 2(LD + K)\gamma$. This implies that $\Delta(X, X', \mathbf{w})$ is a $(LD + K)\gamma$ -sub-Gaussian random variable. For any finite subset $\mathcal{Z}' \subset \mathcal{Z}$ of our convex set, we can thus apply the maximal inequality to obtain a uniform bound over all potential values of the variational term \mathbf{w} :

$$\mathbb{E}_{X, X' \sim \mathcal{D}^n} \left[\max_{\mathbf{w} \in \mathcal{Z}'} \Delta(X, X', \mathbf{w}) \right] \leq \sqrt{2}(LD + K)\gamma \log(|\mathcal{Z}'|).$$

We can thus apply Jensen's inequality and (2) to bound the generalization gap when \mathbf{w} is restricted to the subset \mathcal{Z}' :

$$\sqrt{2}(LD + K)\gamma \log(|\mathcal{Z}'|) \geq \mathbb{E}_{X \sim \mathcal{D}^n} \left[\max_{\mathbf{w} \in \mathcal{Z}'} \langle F(A(X)), A(X) - \mathbf{w} \rangle \right] - \mathbb{E}_{X \sim \mathcal{D}^n} [\text{Err}_{\widehat{F}}(A(X))]. \quad (3)$$

We now choose the subset \mathcal{Z}' to cover the domain \mathcal{Z} , proceeding with a one-step discretization for ease of exposition. Let \mathcal{Z}_δ be the smallest δ -covering of the set \mathcal{Z} in its norm $\|\cdot\|$; i.e., for each

$\mathbf{z} \in \mathcal{Z}$, there exists $\mathbf{z}_\delta \in \mathcal{Z}_\delta$ such that $\|\mathbf{z} - \mathbf{z}_\delta\| \leq \delta$. Fixing $\delta > 0$, we can bound the generalization error by

$$\begin{aligned} \mathbb{E}_{X \sim \mathcal{D}^n} [\text{Err}_F(A(X)) - \text{Err}_{\hat{F}}(A(X))] &= \mathbb{E}_{X \sim \mathcal{D}^n} \left[\max_{\mathbf{w} \in \mathcal{Z}_\delta} \langle F(A(X)), A(X) - \mathbf{w} \rangle \right] - \mathbb{E}_{X \sim \mathcal{D}^n} [\text{Err}_{\hat{F}}(A(X))] \\ &+ \mathbb{E}_{X \sim \mathcal{D}^n} [\text{Err}_F(A(X))] - \mathbb{E}_{X \sim \mathcal{D}^n} \left[\max_{\mathbf{w} \in \mathcal{Z}_\delta} \langle F(A(X)), A(X) - \mathbf{w} \rangle \right]. \end{aligned}$$

We can upper bound the first two summands by (3) to be at most

$$\mathbb{E}_{X \sim \mathcal{D}^n} \left[\max_{\mathbf{w} \in \mathcal{Z}_\Delta} \langle F(A(X)), A(X) - \mathbf{w} \rangle \right] - \mathbb{E}_{X \sim \mathcal{D}^n} [\text{Err}_{\hat{F}}(A(X))] \leq \sqrt{2}(LD + K)\gamma \log(|\mathcal{Z}_\delta|).$$

As \mathcal{Z}_δ is a δ -cover, the next two summands are at most

$$\mathbb{E}_{X \sim \mathcal{D}^n} [\text{Err}_F(A(X))] - \mathbb{E}_{X \sim \mathcal{D}^n} \left[\max_{\mathbf{w} \in \mathcal{Z}_{2-M}} \langle F(A(X)), A(X) - \mathbf{w} \rangle \right] \leq K\delta.$$

□

Due to the linear structure of the inner product in the gap function, it is possible to control the generalization error by taking a union bound only over the extremal points of the domain. This means that, in certain settings, we can improve on the rates of Theorem 3.1. For example, let us consider the $(d + 1)$ -dimensional probability simplex Δ_d equipped with the L1 norm.

Theorem 3.2. *Let (F, Δ_d) be a smooth strongly monotone variational inequality (Assumptions 2.1, 2.3, 2.5). The expected generalization gap for a γ -stable algorithm $A: \mathcal{X}^n \rightarrow \Delta_d$ is:*

$$\mathbb{E}_{X \sim \mathcal{D}^n} [\text{Err}_F(A(X)) - \text{Err}_{\hat{F}}(A(X))] \in O((L + K)\gamma \log d).$$

Proof. Because linear functions are maximized on the vertices of a simplex, we have that

$$\min_{w \in \Delta_d} \langle F(A(X)), w \rangle = \min_{i \in [d+1]} \langle F(A(X)), e_i \rangle,$$

where e_i is the i th standard basis vector. It thus suffices to only vary w over the set $\mathcal{Z}' = \{e_i\}_{i \in [d+1]}$. The claim then follows by (3). □

We can, for example, plug the $\gamma = \frac{2K}{n(2\mu - \eta L^2)}$ stability of the gradient descent algorithm (Lemma 2.8) into Theorem 3.2 to obtain a fast $O(1/n)$ generalization error bound in strongly monotone VIs.

Corollary 3.3. *Let (F, Δ_d) be a smooth strongly monotone variational inequality (Assumptions 2.1, 2.3, 2.5). The expected generalization error of the gradient descent algorithm $A: \mathcal{X}^n \rightarrow \Delta_d$ with sufficiently small learning rate $\eta \rightarrow 0$ and large time horizon $T \rightarrow \infty$ is:*

$$\mathbb{E}_{X \sim \mathcal{D}^n} [\text{Err}_F(A(X))] \leq \tilde{O}\left(\frac{K \cdot (L + K)}{n\mu}\right),$$

where \tilde{O} hides logarithmic factors.

By Lemma 2.10, a fast $O(1/n)$ rate can be similarly obtained for extragradient descent.

4 Fast Potential Gap Bounds for Games

We now turn to the case of variational inequalities that describe multi-player games, where we assume that the operator F is conservative and the domain has the product structure $\mathcal{Z} = \mathcal{Z}_1, \dots, \mathcal{Z}_k$. Throughout this section, we let F_i denote the restriction of the operator F to player i 's strategy space \mathcal{Z}_i , and use μ_i and L_i to denote player i 's strong monotonicity constant and the smoothness constant respectively; note that $\mu_i \geq \mu$ and $L_i \leq L$.

Another method of managing the instability of the gap function is to avoid bounding the generalization error of the gap function itself and only seek to bound the *potential gap*. In many applications, such as finding Nash equilibrium in multi-player games, only the latter is of interest and the gap function serves only as a means to upper bound the potential gap. In these cases, we can define a surrogate gap function that replaces the unstable term $\arg \max_{\mathbf{z}^* \in \mathcal{Z}} \langle F(\mathbf{z}), \mathbf{z} - \mathbf{z}^* \rangle$ with

$$w^*(\mathbf{z}) = \left[\arg \min_{\mathbf{z}_i^* \in \mathcal{Z}_i} f_i(\mathbf{z}_i^*, \mathbf{z}_{-i}) \right]_{i \in [k]}. \quad (4)$$

We denote this weak form of the gap function as

$$\text{WeakGap}_F(\mathbf{z}) \triangleq \langle F(\mathbf{z}), \mathbf{z} - w^*(\mathbf{z}) \rangle,$$

and note that it also upper bounds the potential gap by $\text{Err}_F(\mathbf{z}) \leq \text{WeakGap}_F(\mathbf{z})$. Unlike the gap function, this weaker gap function is stable due to the stability of the term $w^*(\mathbf{z})$, as noted by Lin et al. (2020) in the min-max settings. We can also observe that the empirical form of this weak gap function, $\widehat{\text{WeakGap}}_X(\mathbf{z}) = \langle \widehat{F}(\mathbf{z}), \mathbf{z} - w^*(\mathbf{z}) \rangle$, is linear in the dataset X since the term $w^*(\mathbf{z})$ is fixed. This allows us to reduce the problem of upper bounding the potential gap to the familiar problem of bounding the generalization error of an expected loss function.

Theorem 4.1. *Let (F, \mathcal{Z}) be a smooth strongly monotone variational inequality (Assumptions 2.1, 2.3, 2.5) corresponding to a k -player game. The expected generalization gap for a γ -stable algorithm $A: \mathcal{X}^n \rightarrow \mathcal{Z}$ is*

$$\mathbb{E}_{X \sim \mathcal{D}^n} [\text{WeakGap}_F(A(X)) - \text{WeakGap}_{\widehat{F}}(A(X))] \in O\left(\gamma(2DL + K \sum_{i \in [k]} \frac{L_i}{\mu_i})\right).$$

Proof. The usual symmetrization argument implies that expected generalization gap is bounded by

$$\begin{aligned} \mathbb{E}_{X \sim \mathcal{D}^n} [\text{WeakGap}_F(A(X))] &= \frac{1}{n} \sum_{j=1}^n \mathbb{E}_{\substack{X \sim \mathcal{D}^n \\ X' \sim \mathcal{D}^n}} \left[\langle \widehat{F}_{X_j'}(A(X)), A(X) - w^*(A(X)) \rangle \right] \\ &= \frac{1}{n} \sum_{j=1}^n \mathbb{E}_{\substack{X \sim \mathcal{D}^n \\ X' \sim \mathcal{D}^n}} \left[\langle \widehat{F}_{X_j}(A(X^{(j)})), A(X^{(j)}) - w^*(A(X^{(j)})) \rangle \right] \\ &\leq \Delta + \mathbb{E}_{X \sim \mathcal{D}^n} [\text{WeakGap}_{\widehat{F}_X}(A(X))], \end{aligned}$$

where Δ is the quantity

$$\Delta = \left\langle \widehat{F}_{X_j}(A(X^{(j)})), A(X^{(j)}) - w^*(A(X^{(j)})) \right\rangle - \text{WeakGap}_{\widehat{F}_X}(A(X)),$$

that we can decompose into the terms

$$(1) = \left\langle \widehat{F}_{X_j}(A(X^{(j)})), A(X^{(j)}) \right\rangle - \left\langle \widehat{F}_{X_j}(A(X)), A(X) \right\rangle,$$

and

$$(2) = \left\langle \widehat{F}_{X_j}(A(X)), w^*(A(X)) \right\rangle - \left\langle \widehat{F}_{X_j}(A(X^{(j)})), w^*(A(X^{(j)})) \right\rangle.$$

We can control first term using Hölder's inequality and the smoothness of \widehat{F}_{X_j} to obtain

$$\begin{aligned} (1) &= \left\langle \widehat{F}_{X_i}(A(X^{(i)})) - \widehat{F}_{X_i}(A(X)), A(X^{(i)}) \right\rangle + \left\langle \widehat{F}_{X_i}(A(X)), A(X^{(i)}) - A(X) \right\rangle \\ &\leq DL\gamma + K\gamma. \end{aligned} \quad (5)$$

We can similarly bound the second term by

$$\begin{aligned} (2) &= \left\langle \widehat{F}_{X_j}(A(X)) - \widehat{F}_{X_j}(A(X^{(j)})), w^*(A(X^{(j)})) \right\rangle + \left\langle \widehat{F}_{X_j}(A(X)), w^*(A(X)) - w^*(A(X^{(j)})) \right\rangle \\ &\leq DL\gamma + K \left\| w^*(A(X)) - w^*(A(X^{(j)})) \right\|. \end{aligned} \quad (6)$$

We now bound $\left\| w^*(A(X)) - w^*(A(X^{(j)})) \right\|$ in terms of $\left\| A(X) - A(X^{(j)}) \right\|$ using the technique of Lin et al. (2020). This is the step that was not possible with the original gap function. We proceed by bounding each player's contribution to this norm separately. Due to the optimality of $w^*(z)_i$, the restriction of $w^*(z)$ to the i th player's strategy space \mathcal{Z}_i , the first-order optimality condition of the strongly convex function f_i implies that, for all $\mathbf{x}, \mathbf{x}' \in \mathcal{Z}_i$,

$$\langle F(w^*(z)_i, z_{-i})_i, w^*(z)_i - \mathbf{x} \rangle \leq 0 \quad \text{and} \quad \langle F(w^*(z')_i, z'_{-i})_i, w^*(z')_i - \mathbf{x}' \rangle \leq 0.$$

Choosing $\mathbf{x} = w^*(z')_i$ and $\mathbf{x}' = w^*(z)_i$ and summing the two inequalities gives

$$\langle F(w^*(z)_i, z_{-i})_i - F(w^*(z')_i, z'_{-i})_i, w^*(z)_i - w^*(z')_i \rangle \leq 0.$$

In concert with the strong monotonicity of F_i , which implies

$$\langle F(w^*(z)_i, z_{-i})_i - F(w^*(z')_i, z_{-i})_i, w^*(z)_i - w^*(z')_i \rangle - \mu_i \left\| w^*(z)_i - w^*(z')_i \right\|^2 \geq 0,$$

we have

$$\begin{aligned} \mu_i \left\| w^*(z)_i - w^*(z')_i \right\|^2 &\leq \langle F(w^*(z')_i, z'_{-i})_i - F(w^*(z')_i, z_{-i})_i, w^*(z)_i - w^*(z')_i \rangle \\ &\leq L_i \left\| z'_{-i} - z_{-i} \right\| \left\| w^*(z)_i - w^*(z')_i \right\|. \end{aligned} \quad (7)$$

Choosing $\mathbf{z} = A(X)$ and $\mathbf{z}' = A(X^{(j)})$ and substituting the inequality into (6) gives

$$(2) \leq DL\gamma + K \sum_{i=1}^k \frac{L_i}{\mu_i} \left\| A(X)_i - A(X^{(j)})_i \right\|. \quad (8)$$

Summing (8) with (5) completes the proof. \square

In this approach, we reduce the problem of bounding the potential gap to bounding the generalization error of an expected loss. One implication is that we can directly apply—rather than reinventing—the many technical tools developed for the stability analysis of convex minimization. For example, we can improve Theorem 4.1 to the following high-probability guarantee due to Klochkov and Zhivotovskiy (2021).

Theorem 4.2. *Let (F, \mathcal{Z}) be a smooth strongly monotone variational inequality (Assumptions 2.1, 2.3, 2.5) corresponding to a k -player game. The expected generalization gap for a γ -stable algorithm $A: \mathcal{X}^n \rightarrow \mathcal{Z}$ when training error is small; i.e., $\text{WeakGap}_{\widehat{F}}(A(X)) \leq \varepsilon$, is, with probability at least $1 - \delta$,*

$$\text{WeakGap}_F(A(X)) \leq 2\varepsilon + O\left(\log\left(\frac{1}{\delta}\right)\left((DL + K \sum_{i \in [k]} \frac{L_i}{\mu_i})\gamma \log n + \frac{(DL + K \sum_{i \in [k]} \frac{L_i}{\mu_i})^2}{n}\right)\right).$$

This provides a fast-rate high-probability potential gap bound of order $O(1/n)$ for gradient descent in games.

Corollary 4.3. *Let (F, \mathcal{Z}) be a smooth strongly monotone variational inequality (Assumptions 2.1, 2.3, 2.5) corresponding to a k -player game. Running the gradient descent algorithm $A: \mathcal{X}^n \rightarrow \Delta_d$ with sufficiently small learning rate $\eta \rightarrow 0$ and large time horizon $T \rightarrow \infty$ results in an iterate $A(X)$ that with probability at least $1 - \delta$ has a potential gap bounded by*

$$\text{Err}_F(A(X)) \leq \widetilde{O}\left(\frac{\log(1/\delta)}{n}\left(\frac{K}{\mu}(DL + K \sum_{i \in [k]} \frac{L_i}{\mu_i}) + (DL + K \sum_{i \in [k]} \frac{L_i}{\mu_i})^2\right)\right),$$

where \widetilde{O} hides logarithmic factors.

Once again, by Lemma 2.10, a fast $O(1/n)$ rate can be similarly obtained for extragradient descent.

Theorem 4.2 follows directly from proving that a Bernstein noise condition (Koltchinskii, 2006; Klochkov and Zhivotovskiy, 2021) is satisfied by the weak gap function.

Definition 4.4 (Generalized Bernstein Noise Condition for Minimization). *We say that an algorithm $A: \mathcal{X}_n \times \mathcal{X} \rightarrow [0, M]$ satisfies the generalized Bernstein assumption for some parameter B if, for any $B > 0$, $\mathbf{z} \in \mathcal{Z}$, there is $\mathbf{z}^* \in \mathcal{Z}^*$ such that:*

$$\mathbb{E}_{x' \sim \mathcal{D}} [(f(\mathbf{z}, x') - f(\mathbf{z}^*, x'))^2] \leq B \left(\mathbb{E}_{x' \sim \mathcal{D}} [f(\mathbf{z}, x')] - \mathbb{E}_{x' \sim \mathcal{D}} [f(\mathbf{z}^*, x')] \right).$$

Indeed, this noise condition is satisfied in VIs when strong monotonicity holds.

Fact 4.5. *The weak gap function $\langle F(\mathbf{z}), \mathbf{z} - w^*(\mathbf{z}) \rangle$ of a strongly monotone VI (F, \mathcal{Z}) satisfying the assumptions of Theorem 4.2 fulfills the Bernstein condition with $B = (LD + K(1 + \sum_{i \in [k]} \frac{L_i}{\mu_i}))^2$.*

Proof. We want to upper bound the variance of the weak gap function:

$$\begin{aligned} \text{GapVar} &\triangleq \mathbb{E}[(\langle \widehat{F}(\mathbf{z}), \mathbf{z} - w^*(\mathbf{z}) \rangle - \langle \widehat{F}(\mathbf{z}^*, \mathbf{z}), \mathbf{z}^* - w^*(\mathbf{z}) \rangle)^2] \\ &= \mathbb{E}\left[\left(\langle \widehat{F}(\mathbf{z}) - \widehat{F}(\mathbf{z}^*, \mathbf{z}), \mathbf{z} - w^*(\mathbf{z}) \rangle - \langle \widehat{F}(\mathbf{z}^*, \mathbf{z}), \mathbf{z}^* - \mathbf{z} - w^*(\mathbf{z}^*) + w^*(\mathbf{z}) \rangle\right)^2\right]. \end{aligned}$$

To control the first summand, we use Hölder's inequality to bound

$$\left\langle \widehat{F}(\mathbf{z}) - \widehat{F}(\mathbf{z}^*, \mathbf{z}), \mathbf{z} - w^*(\mathbf{z}) \right\rangle \leq \left\| \widehat{F}(\mathbf{z}) - \widehat{F}(\mathbf{z}^*, \mathbf{z}) \right\| \|\mathbf{z} - w^*(\mathbf{z})\| \leq LD \|\mathbf{z} - \mathbf{z}^*\|.$$

To control the second summand, we can again use Hölder's inequality to bound

$$\begin{aligned} \left\langle \widehat{F}(\mathbf{z}^*, \mathbf{z}), \mathbf{z}^* - \mathbf{z} - w^*(\mathbf{z}^*) + w^*(\mathbf{z}) \right\rangle &\leq K(\|\mathbf{z}^* - \mathbf{z}\| + \|w^*(\mathbf{z}) - w^*(\mathbf{z}^*)\|) \\ &\leq K\left(1 + \sum_{i \in [k]} \frac{L_i}{\mu_i}\right) \|\mathbf{z}^* - \mathbf{z}\|, \end{aligned}$$

where the last step bounds the stability of the variational term surrogate w^* in the same fashion as (7). Summing these bounds, we have as desired

$$\begin{aligned} \text{GapVar} &\leq \mathbb{E} \left[(LD + K(1 + \sum_{i \in [k]} \frac{L_i}{\mu_i}))^2 \|\mathbf{z}^* - \mathbf{z}\|^2 \right] \\ &\leq \mathbb{E} \left[(LD + K(1 + \sum_{i \in [k]} \frac{L_i}{\mu_i}))^2 \langle F(\mathbf{z}), \mathbf{z} - \mathbf{z}^* \rangle \right] \\ &\leq \mathbb{E} \left[(LD + K(1 + \sum_{i \in [k]} \frac{L_i}{\mu_i}))^2 \langle F(\mathbf{z}), \mathbf{z} - w^*(\mathbf{z}) \rangle \right]. \end{aligned}$$

□

The high-probability bound then follows directly from Klochkov and Zhivotovskiy (2021) which provides a high-probability generalization error bound of order $O((\gamma \log(n) + B/n) \log(1/\delta))$ when the Bernstein condition holds with parameter B . We defer the complete proof of Theorem 4.2 to the Appendix.

5 Discussion

This paper aims to provide an accessible exposition for deriving fast-rate generalization bounds in strongly monotone variational inequalities. These derivations, which are direct extensions of stability arguments in convex optimization settings, demonstrate that the generality of working with VIs comes essentially “for free” when it comes to stability-based analyses of first-order methods. There remain, however, gaps in our understanding of learning in strongly monotone VIs. First, there seems to be an inherent hardness to minimizing gap functions with fast rates. Our results on bounding the gap function relied on either assuming a small domain (Theorem 3.1), which can introduce dimension dependence, or assuming domain structure (Theorem 3.2). Second, the generalization rates we derive are likely not optimal in their dependence on condition numbers and especially their dependence on D , the diameter of the domain. We hypothesize that the latter dependence on D may be inevitable when working with gap functions, but it is unclear if this should also be the case for potential gaps.

Acknowledgments

This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE 2146752, and by the European Union (ERC-2022-SYG-OCEAN-101071601). Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the National Science Foundation, the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authorities can be held responsible for them. The authors acknowledge support from the Swiss National Science Foundation (SNSF), grants 199740 and 214441.

References

W. Azizian, I. Mitliagkas, S. Lacoste-Julien, and G. Gidel. A tight and unified analysis of gradient-based methods for a whole spectrum of differentiable games. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 2863–2873, 2020.

- A. Beck and M. Teboulle. Mirror descent and nonlinear projected subgradient methods for convex optimization. *Operations Research Letters*, 31(3):167–175, 2003.
- R. I. Bot, E. R. Csetnek, and D.-K. Nguyen. Fast OGDA in continuous and discrete time. 2022.
- O. Bousquet and A. Elisseeff. Stability and generalization. *Journal of Machine Learning Research (JMLR)*, 2:499–526, 2002.
- Y. Cai, A. Oikonomou, and W. Zheng. Tight last-iterate convergence of the extragradient method for constrained monotone variational inequalities. 2022.
- Z. Charles and D. Papailiopoulos. Stability and generalization of learning algorithms that converge to global optima. In *Proceedings of the International Conference on Machine Learning (ICML)*, volume 80 of *Proceedings of Machine Learning Research*, pages 745–754. PMLR, 2018.
- T. Chavdarova, M. I. Jordan, and M. Zampetakis. Last-iterate convergence of saddle point optimizers via high-resolution differential equations. In *Minimax Theory and its Applications*, 2023.
- R. W. Cottle and G. B. Dantzig. Complementary pivot theory of mathematical programming. *Linear Algebra and its Applications*, 1(1):103–125, 1968.
- C. Daskalakis, A. Ilyas, V. Syrgkanis, and H. Zeng. Training GANs with optimism. In *International Conference on Learning Representations (ICLR)*, 2018.
- L. Devroye and T. Wagner. Distribution-free inequalities for the deleted and holdout error estimates. *Information Theory, IEEE Transactions on*, 25:202 – 207, 1979.
- F. Facchinei and J.-S. Pang. *Finite-Dimensional Variational Inequalities and Complementarity Problems Vol I*. Springer-Verlag, 2003.
- F. Farnia and A. Ozdaglar. Train simultaneously, generalize better: Stability of gradient-based minimax learners. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2021.
- V. Feldman and J. Vondrak. High probability generalization bounds for uniformly stable algorithms with nearly optimal rate, 2019.
- A. Galeotti, S. Goyal, M. O. Jackson, F. Vega-Redondo, and L. Yariv. Network games. *The Review of Economic Studies*, 2010.
- G. Gidel, H. Berard, P. Vincent, and S. Lacoste-Julien. A variational inequality perspective on generative adversarial nets. In *International Conference on Learning Representations (ICLR)*, 2019.
- N. Golowich, S. Pattathil, C. Daskalakis, and A. Ozdaglar. Last iterate is slower than averaged iterate in smooth convex-concave saddle point problems. In *Conference on Learning Theory (COLT)*, pages 1758–1784, 2020.
- E. Gorbunov, H. Berard, G. Gidel, and N. Loizou. Stochastic extragradient: General analysis and improved rates. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2022a.

- E. Gorbunov, N. Loizou, and G. Gidel. Extragradient method: $\mathcal{O}(1/K)$ last-iterate convergence for monotone variational inequalities and connections with cocoercivity. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2022b.
- M. Hardt, B. Recht, and Y. Singer. Train faster, generalize better: Stability of stochastic gradient descent. 2016.
- Y.-G. Hsieh, F. Iutzeler, J. Malick, and P. Mertikopoulos. Explore aggressively, update conservatively: Stochastic extragradient methods with variable stepsize scaling. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- A. Juditsky, A. Nemirovski, and C. Tauvel. Solving variational inequalities with stochastic mirror-prox algorithm. *Stochastic Systems*, 2011. Publisher: INFORMS.
- Y. Klochkov and N. Zhivotovskiy. Stability and deviation optimal risk bounds with convergence rate $\mathcal{O}(1/n)$. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- V. Koltchinskii. Local Rademacher complexities and oracle inequalities in risk minimization. *The Annals of Statistics*, 34, 2006.
- G. M. Korpelevich. The extragradient method for finding saddle points and other problems. *Matecon*, 1976.
- V. Krishna. *Auction Theory*. Elsevier Science, 2009.
- I. Kuzborskij and C. H. Lampert. Data-dependent stability of stochastic gradient descent. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2017.
- Y. Lei and Y. Ying. Fine-grained analysis of stability and generalization for stochastic gradient descent. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2020.
- S. Li and Y. Liu. High probability generalization bounds with fast rates for minimax problems. In *International Conference on Learning Representations (ICLR)*, 2022.
- T. Lin, C. Jin, and M. I. Jordan. On gradient descent ascent for nonconvex-concave minimax problems. In *Proceedings of the International Conference on Machine Learning (ICML)*, volume 119 of *Proceedings of Machine Learning Research*, pages 6083–6093. PMLR, 2020.
- N. Loizou, H. Berard, G. Gidel, I. Mitliagkas, and S. Lacoste-Julien. Stochastic gradient descent-ascent and consensus optimization for smooth games: Convergence analysis under expected co-coercivity. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- R. Lowe, Y. Wu, A. Tamar, J. Harb, P. Abbeel, and I. Mordatch. Multi-agent actor-critic for mixed cooperative-competitive environments. *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.
- J.-E. Martínez-Legaz. Generalized conjugation and related topics, in generalized convexity and fractional programming with economic applications. *Generalized Convexity and Fractional Programming with Economic Applications (Pisa, 1988)*, page 168–197, 1990.
- R. P. McAfee and J. McMillan. Auctions and bidding. *Journal of Economic Literature*, 25:699–738, 1987.

- P. Milgrom. *Putting Auction Theory to Work*. Churchill Lectures in Economics. Cambridge University Press, 2004.
- A. Mokhtari, A. Ozdaglar, and S. Pattathil. A unified analysis of extra-gradient and optimistic gradient methods for saddle point problems: Proximal point approach. 2019.
- A. Mokhtari, A. Ozdaglar, and S. Pattathil. Convergence rate of $\mathcal{O}(1/k)$ for optimistic gradient and extra-gradient methods in smooth convex-concave saddle point problems. In *SIAM Journal on Optimization*, 2020.
- S. Mukherjee, P. Niyogi, T. Poggio, and R. Rifkin. Learning theory: Stability is sufficient for generalization and necessary and sufficient for consistency of empirical risk minimization. *Adv. Comput. Math.*, 25:161–193, 2006.
- A. Nagurney and D. Zhang. Projected dynamical systems and variational inequalities with applications. In *International Series in Operations Research and Management Science (ISOR)*, volume 2, 1995.
- A. Nemirovski. Prox-method with rate of convergence $o(1/t)$ for variational inequalities with Lipschitz continuous monotone operators and smooth convex-concave saddle point problems. *SIAM Journal on Optimization*, 15(1):229–251, 2004.
- A. Nemirovski and D. Yudin. *Problem Complexity and Method Efficiency in Optimization*, volume 1. Wiley, New York, 1983.
- F. Parise and A. E. Ozdaglar. A variational inequality framework for network games: Existence, uniqueness, convergence and sensitivity analysis. *Games and Economic Behavior*, 114:47–82, 2019.
- L. D. Popov. A modification of the Arrow–Hurwicz method for search of saddle points. *Mathematical Notes of the Academy of Sciences of the USSR*, 28(5):845–848, 1980.
- J. G. Riley. Expected revenue from open and sealed bid auctions. *Journal of Economic Perspectives*, pages 41–50, 1989.
- W. H. Rogers and T. J. Wagner. A finite sample distribution-free performance bound for local discrimination rules. *Annals of Statistics*, 6:506–514, 1978.
- G. Scutari, D. P. Palomar, F. Facchinei, and J.-s. Pang. Convex optimization, game theory, and variational inequality theory. *IEEE Signal Processing Magazine*, 27(3):35–49, 2010.
- S. Shalev-Shwartz, O. Shamir, N. Srebro, and K. Sridharan. Learnability, stability and uniform convergence. *J. Mach. Learn. Res.*, 11:2635–2670, 2010.
- B. A. M. Sidahmed and T. Chavdarova. Variational inequality methods for multi-agent reinforcement learning: Performance and stability gains. 2024.
- G. Stampacchia. Formes bilinéaires coercitives sur les ensembles convexes. *Académie des Sciences de Paris*, 258:4413–4416, 1964.
- P. Tseng. On linear convergence of iterative methods for the variational inequality problem. *Journal of Computational and Applied Mathematics*, 1995.

- O. Vinyals, T. Ewalds, S. Bartunov, P. Georgiev, A. S. Vezhnevets, M. Yeo, A. Makhzani, H. Küttler, J. Agapiou, J. Schrittwieser, J. Quan, S. Gaffney, S. Petersen, K. Simonyan, T. Schaul, H. van Hasselt, D. Silver, T. Lillicrap, K. Calderone, P. Keet, A. Brunasso, D. Lawrence, A. Ekermo, J. Repp, and R. Tsing. StarCraft II: A new challenge for reinforcement learning, 2017.
- T. Yang, M. I. Jordan, and T. Chavdarova. Solving constrained variational inequalities via a first-order interior point-based method. In *International Conference on Learning Representations (ICLR)*, 2023.
- J. Zhang, M. Hong, M. Wang, and S. Zhang. Generalization bounds for stochastic saddle point problems, 2020.

A Additional Preliminaries

A.1 VI Methods

Algorithms 1 and 2 describe the (GD) and (EG) methods, respectively.

Algorithm 1: Gradient Descent

Input : $z_0 \in \mathcal{Z}^o$, step size $\eta > 0$, number of iterations $T > 0$
for $t = 1, \dots, T$ **do**
 | $z_t = z_{t-1} - \eta F(z_{t-1})$.
end
Output: z_t , and the average iterate
 $\bar{z}_t = \frac{1}{t} \sum_{s=1}^t z_s$

Algorithm 2: Extragradient

Input : $z_0 \in \mathcal{Z}^o$, step size $\eta > 0$, iterations $T > 0$
for $t = 1, \dots, T$ **do**
 | $z_{t-\frac{1}{2}} = z_{t-1} - \eta F(z_{t-1})$
 | $z_t = z_{t-1} - \eta F(z_{t-\frac{1}{2}})$
end
Output: z_t , and the average $\bar{z}_t = \frac{1}{t} \sum_{s=1}^t z_s$

A.2 Additional definitions

Contractive operators are defined as follows.

Definition A.1 (Contractiveness and non-expansiveness). *We say a mapping $G: \mathcal{Z} \rightarrow \mathbb{R}^d$, with $\mathcal{Z} \subseteq \mathbb{R}^d$, is non-expansive or c -contractive on \mathcal{Z} if there exists $c \in (0, 1]$ or $c \in (0, 1)$, respectively, such that:*

$$\|G(z) - G(z')\| \leq c \|z - z'\|, \quad \forall z, z' \in \mathcal{Z}.$$

B Proofs

B.1 Growth lemma and proof of Lemma 2.8

The growth lemma formalizes the notion of running a contractive iterative algorithm (as per Def. A.1) on two different but similar datasets.

Lemma B.1 (Generalized growth lemma (Farnia and Ozdaglar, 2021) Lemma 2). *Consider two sequences of updates, $\tilde{P}_1, \dots, \tilde{P}_T$ and $\tilde{P}'_1, \dots, \tilde{P}'_T$, where each update is a mapping of form $\tilde{P}_i: \mathcal{Z} \rightarrow \mathcal{Z}$. Let $P: \mathcal{Z} \rightarrow \mathcal{Z}$ be some mapping, such that $P + \tilde{P}_1, \dots, P + \tilde{P}_T$ and $P + \tilde{P}'_1, \dots, P + \tilde{P}'_T$ are two sequences of ξ -contractive mappings that each take \mathcal{Z} to \mathcal{Z} . Fix some $z \in \mathcal{Z}$ and define $z_1, \dots, z_T \in \mathcal{X}$ to be the trajectory unrolled by applying the updates $P + \tilde{P}_1, \dots, P + \tilde{P}_T$ in sequence, starting at z . Analogously define $z'_1, \dots, z'_T \in \mathcal{X}$ with respect to $P + \tilde{P}'_1, \dots, P + \tilde{P}'_T$. For timesteps*

t where $\tilde{P}_t = \tilde{P}'_t$, we have $\|z_{t+1} - z'_{t+1}\| \leq \xi \|z_t - z'_t\|$. For all timesteps t , the following holds for every constant $r > 0$:

$$\|z_{t+1} - z'_{t+1}\| \leq \xi \|z_t - z'_t\| + \sup_{z \in \mathcal{Z}} \|rz - \tilde{P}_t(z)\| + \sup_{z \in \mathcal{Z}} \|rz - \tilde{P}'_t(z)\|.$$

Proof of Lemma B.1. The first claim holds by definition of contractiveness. The second claim holds by fixing any $r > 0$ and observing that

$$\begin{aligned} \|z_{t+1} - z'_{t+1}\| &= \left\| P(z_t) - P(z'_t) + \tilde{P}_t(z_t) - \tilde{P}'_t(z'_t) \right\| \\ &= \left\| P(z_t) - P(z'_t) + \tilde{P}_t(z_t) - \tilde{P}_t(z'_t) + \tilde{P}_t(z'_t) - \tilde{P}'_t(z'_t) \right\| \\ &\leq \left\| P(z_t) + \tilde{P}_t(z_t) - P(z'_t) - \tilde{P}_t(z'_t) \right\| + \left\| \tilde{P}_t(z'_t) - \tilde{P}'_t(z'_t) \right\| \\ &\leq \xi \|z_t - z'_t\| + \left\| \tilde{P}_t(z'_t) - rz'_t \right\| + \left\| rz'_t - \tilde{P}'_t(z'_t) \right\| \\ &\leq \xi \|z_t - z'_t\| + \sup_{z \in \mathcal{Z}} \left\| \tilde{P}_t(z) - rz \right\| + \sup_{z \in \mathcal{Z}} \left\| rz - \tilde{P}'_t(z) \right\|. \end{aligned}$$

The first equality holds by definition of the iterate $z_{t+1} = \left\| P + \tilde{P}_t(z_t) \right\|$ and likewise for z'_{t+1} . The first inequality is an application of the triangle inequality. The second inequality uses the ξ -contractiveness of $P + \tilde{P}_t$. The third inequality relaxes z'_t to be a supremum over the whole domain \mathcal{X} . \square

For completeness, we retrace below a standard proof of the contractiveness of gradient descent on strongly convex functions but in the more general setting of strongly monotone VIs—observing that the proof goes through without modification.

Lemma B.2 (GD contractiveness). *Let $F: \mathcal{Z} \rightarrow \mathbb{R}^d$ be a μ -strongly monotone (Assumption 2.3) and L -Lipschitz (Assumption 2.1) operator. For any learning rate $\gamma > 0$:*

$$\|(z - \gamma F(z)) - (z' - \gamma F(z'))\|_2^2 \leq (1 - 2\gamma\mu + \gamma^2 L^2) \|z - z'\|_2^2, \quad \forall z, z' \in \mathcal{Z}.$$

Moreover, if the learning rate satisfies $0 < \gamma < \frac{2\mu}{L^2}$:

$$\|(z - \gamma F(z)) - (z' - \gamma F(z'))\|_2^2 \leq \left(1 - \frac{\mu^2}{L^2}\right) \|z - z'\|_2^2, \quad \forall z, z' \in \mathcal{Z}.$$

Proof of Lemma B.2. Developing the left-hand-side of our first claim:

$$\begin{aligned} \|(z - \alpha F(z)) - (z' - \alpha F(z'))\|_2^2 &= \|z - z' - \alpha(F(z) - F(z'))\|_2^2 \\ &= \|z - z'\|_2^2 - 2\alpha \underbrace{\langle F(z) - F(z'), z - z' \rangle}_{\geq \mu \|z - z'\|_2^2} + \alpha^2 \underbrace{\|F(z) - F(z')\|_2^2}_{\leq L^2 \|z - z'\|_2^2} \\ &\leq (1 - 2\alpha\mu + \alpha^2 L^2) \|z - z'\|_2^2. \end{aligned}$$

The second equality just expands the square, while the first inequality uses the μ -strong monotonicity of F to bound $\langle F(z) - F(z'), z - z' \rangle \geq \mu \|z - z'\|_2^2$ and the L -Lipschitzness of F to bound $\|F(z) - F(z')\|_2^2 \leq L^2 \|z - z'\|_2^2$. \square

We now proceed to bound the stability of the iterates of gradient descent.

Lemma 2.8 (GD stability). *Consider a smooth strongly monotone variational inequality (F, \mathcal{Z}) (Assumption 2.1, 2.3, 2.5). Given neighboring datasets $X \in \mathcal{X}_n$ and $X' \in \mathcal{X}_n$, let \mathbf{z}_T and \mathbf{z}'_T denote the T th iterates of gradient descent (Algorithm 1) on the empirical operators given by X and X' respectively. For step sizes $0 < \eta < \frac{2\mu}{L^2}$, uniform stability holds with*

$$\|\mathbf{z}_T - \mathbf{z}'_T\| \leq \frac{2K}{n(2\mu - \eta L^2)}.$$

Proof of Lemma 2.8. Let j be the index of \mathcal{X}_n and \mathcal{X}'_n where the datasets disagree: $\xi_j \neq \xi'_j$. By definition of gradient descent, we can express each iterate as:

$$\mathbf{z}_t = \mathbf{z}_{t-1} - \frac{\gamma}{n} \sum_{i=1}^n \Xi(\mathbf{z}_{t-1}, \xi_i), \quad \mathbf{z}'_t = \mathbf{z}'_{t-1} - \frac{\gamma}{n} \sum_{i=1}^n \Xi(\mathbf{z}'_{t-1}, \xi'_i).$$

Thus, we can express $\mathbf{z}_1, \dots, \mathbf{z}_T$ as the trajectory of applying updates $P + \tilde{P}_1, \dots, P + \tilde{P}_T$ where we define:

$$P(\mathbf{z}) = \mathbf{z} - \frac{\gamma}{n} \sum_{i \in [n] \setminus \{j\}} \Xi(\mathbf{z}, \xi_i), \quad \tilde{P}_t(\mathbf{z}) = \frac{\gamma}{n} \Xi(\mathbf{z}, \xi_j).$$

Similarly, we can express $\mathbf{z}'_1, \dots, \mathbf{z}'_T$ in terms of updates $P' + \tilde{P}'_1, \dots, P' + \tilde{P}'_T$ where:

$$P'(\mathbf{z}) = \mathbf{z} - \frac{\gamma}{n} \sum_{i \in [n] \setminus \{j\}} \Xi(\mathbf{z}, \xi'_i), \quad \tilde{P}'_t(\mathbf{z}) = \frac{\gamma}{n} \Xi(\mathbf{z}, \xi'_j).$$

Since $\xi_i = \xi'_i$ for all $i \neq j$, we have that $P = P'$. That is, the updates provided by each dataset disagree only in an additive component that corresponds to the disagreeing datapoint. The growth lemma, Lemma B.1, thus implies:

$$\|\mathbf{z}_{t+1} - \mathbf{z}'_{t+1}\| \leq \xi \|\mathbf{z}_t - \mathbf{z}'_t\| + \|\tilde{P}_t(\mathbf{z}'_t)\| + \|\tilde{P}'_t(\mathbf{z}'_t)\|,$$

where ξ is the contractiveness of $P + \tilde{P}_t$. Since we are in a strongly monotone and smooth setting, Lemma B.2 bounds the contractiveness of gradient descent updates by $\xi \leq (1 - 2\gamma\mu + \gamma^2 L^2)$. Because gradient descent has a bounded learning rate and we assume our operator is bounded, we have $\|\tilde{P}_t(\mathbf{z}'_t)\| \leq \gamma K/n$ and $\|\tilde{P}'_t(\mathbf{z}'_t)\| \leq \gamma K/n$. Thus:

$$\|\mathbf{z}_{t+1} - \mathbf{z}'_{t+1}\| \leq (1 - 2\gamma\mu + \gamma^2 L^2) \|\mathbf{z}_t - \mathbf{z}'_t\| + \frac{2\gamma K}{n}.$$

Recurring through $\mathbf{z}_1, \dots, \mathbf{z}_T$:

$$\|\mathbf{z}_T - \mathbf{z}'_T\| \leq \frac{2\gamma K}{n} \sum_{t=1}^T (1 - 2\gamma\mu + \gamma^2 L^2)^{T-t} \leq \frac{2K}{n(2\mu - \gamma L^2)}$$

where the second inequality is because $\sum_{t=1}^T x^{T-t} = \frac{1-x^T}{1-x}$. □

B.2 Proof of Lemma 2.9 and Lemma 2.10

In this section, we prove the contractiveness of extragradient descent updates in strongly monotone VIs and by extension the stability of its iterates.

Lemma 2.9 (EG contractiveness). *Let $F: \mathcal{Z} \rightarrow \mathbb{R}^d$ be a L -Lipschitz μ -strongly monotone operator (Assumptions 2.1, 2.3). Let $G_{EG}(\mathbf{z}, F)$ denote the EGD update of an iterate \mathbf{z} given operator F , that is $G_{EG} \equiv I - \eta F \circ (I - \eta F)$. For any learning rate $\eta > 0$, and $\forall \mathbf{z}, \mathbf{z}' \in \mathcal{Z}$ we have:*

$$G_{EG}(\mathbf{z}, F) \leq \underbrace{(1 - \eta + \eta^2 L^2 (1 - 2\eta\mu + \eta^2 L^2))}_{\triangleq c_{EG}(\eta)} \|\mathbf{z} - \mathbf{z}'\|_2^2.$$

That is, G_{EG} is contractive for any step size η where $c_{EG}(\eta) < 1$ and $\{\eta_i \mid c_{EG}(\eta_i) < 1\}$ is non empty.

Proof of Lemma 2.9. Developing the left-hand-side of our first claim:

$$\begin{aligned} \|T_{EG}(\mathbf{z}') - T_{EG}(\mathbf{z})\|_2^2 &= \|\mathbf{z}' - \eta F(\mathbf{z}' - \eta F(\mathbf{z}')) - \mathbf{z} + \eta F(\mathbf{z} - \eta F(\mathbf{z}))\|_2^2 \\ &= \|\mathbf{z}' - \mathbf{z}\|_2^2 - \underbrace{2\eta \langle \mathbf{z}' - \mathbf{z}, F(\mathbf{z}' - \eta F(\mathbf{z}')) - F(\mathbf{z} - \eta F(\mathbf{z})) \rangle}_{1^\circ} \\ &\quad + \underbrace{\eta^2 \|F(\mathbf{z}' - \eta F(\mathbf{z}')) - F(\mathbf{z} - \eta F(\mathbf{z}))\|_2^2}_{2^\circ}. \end{aligned} \quad (9)$$

Term 1^o. Using the parallelogram identity—for any two vectors \mathbf{a}, \mathbf{b} we have $2\langle \mathbf{a}, \mathbf{b} \rangle \geq \|\mathbf{a}\|_2^2 + \|\mathbf{b}\|_2^2$ —yields:

$$\begin{aligned} &\langle \mathbf{z}' - \mathbf{z}, F(\mathbf{z}' - \eta F(\mathbf{z}')) - F(\mathbf{z} - \eta F(\mathbf{z})) \rangle \\ &\geq \frac{1}{2} \|\mathbf{z}' - \mathbf{z}\|_2^2 + \frac{1}{2} \|F(\mathbf{z}' - \eta F(\mathbf{z}')) - F(\mathbf{z} - \eta F(\mathbf{z}))\|_2^2 \\ &\geq \frac{1}{2} \|\mathbf{z}' - \mathbf{z}\|_2^2. \end{aligned} \quad (10)$$

Hence,

$$-2\eta \langle \mathbf{z}' - \mathbf{z}, F(\mathbf{z}' - \eta F(\mathbf{z}')) - F(\mathbf{z} - \eta F(\mathbf{z})) \rangle \leq -\eta \|\mathbf{z}' - \mathbf{z}\|_2^2. \quad (11)$$

Term 2^o. For 2^o in (9) we have:

$$\begin{aligned} \eta^2 \|F(\mathbf{z}' - \eta F(\mathbf{z}')) - F(\mathbf{z} - \eta F(\mathbf{z}))\|_2^2 &\leq \eta^2 L^2 \|\mathbf{z}' - \eta F(\mathbf{z}') - \mathbf{z} + \eta F(\mathbf{z})\|_2^2 \\ &\leq \eta^2 L^2 (1 - 2\eta\mu + \eta^2 L^2) \|\mathbf{z}' - \mathbf{z}\|_2^2, \end{aligned} \quad (12)$$

where for the first inequality we used Definition 2.1, and for the latter Lemma B.2. By plugging (11) and (12) in (9) we have:

$$\begin{aligned} \|T_{EG}(\mathbf{z}') - T_{EG}(\mathbf{z})\|_2^2 &= \|\mathbf{z}' - \eta F(\mathbf{z}' - \eta F(\mathbf{z}')) - \mathbf{z} + \eta F(\mathbf{z} - \eta F(\mathbf{z}))\|_2^2 \\ &\leq \underbrace{(1 - \eta)}_{1^\circ} + \underbrace{\eta^2 L^2 (1 - 2\eta\mu + \eta^2 L^2)}_{2^\circ} \|\mathbf{z}' - \mathbf{z}\|_2^2. \end{aligned} \quad (13)$$

Hence, the conclusion follows with:

$$c_{EG}(\eta) = 1 - \eta + \eta^2 L^2 (1 - 2\eta\mu + \eta^2 L^2). \quad (14)$$

Note that for T_{EG} to be non-expansive, it has to hold that $c_{EG} < 1$. \square

Given this contractiveness result, we can apply the growth lemma as usual to bound the stability of extragradient descent.

Lemma 2.10 (EG stability). *Consider a smooth strongly monotone variational inequality (F, \mathcal{Z}) (Assumption 2.1, 2.3, 2.5). Given neighboring datasets $X \in \mathcal{X}_n$ and $X' \in \mathcal{X}'_n$, let \mathbf{z}_T and \mathbf{z}'_T denote the T th iterates of extragradient descent (Algorithm 2) on the empirical operators given by X and X' respectively. For step sizes $\eta > 0$, uniform stability holds with*

$$\|\mathbf{z}_T - \mathbf{z}'_T\| \leq \frac{2K}{n(1 + \eta - \eta L^2(1 - \eta)(1 - 2\eta\mu + \eta^2 L^2))(1 - \eta L - (1 + \eta L))}.$$

Proof of Lemma 2.10. Let j be the index of \mathcal{X}_n and \mathcal{X}'_n where the datasets disagree: $\xi_j \neq \xi'_j$. By definition of extragradient descent, we can express each iterate as

$$\begin{aligned} \mathbf{z}_t &= \mathbf{z}_{t-1} - \frac{\eta}{n} \sum_{i=1}^n \Xi \left(\mathbf{z}_{t-1} - \eta \sum_{i=1}^n \Xi(\mathbf{z}_{t-1}, \xi_i), \xi_i \right), \\ \mathbf{z}'_t &= \mathbf{z}'_{t-1} - \frac{\eta}{n} \sum_{i=1}^n \Xi \left(\mathbf{z}_{t-1} - \eta \sum_{i=1}^n \Xi(\mathbf{z}_{t-1}, \xi'_i), \xi'_i \right). \end{aligned}$$

We can compare this to the outcome of a gradient-descent step:

$$\begin{aligned} \tilde{\mathbf{z}}_t &= \mathbf{z}_{t-1} - \frac{\eta}{n} \sum_{i=1}^n \Xi(\mathbf{z}_{t-1}, \xi_i), \\ \tilde{\mathbf{z}}'_t &= \mathbf{z}'_{t-1} - \frac{\eta}{n} \sum_{i=1}^n \Xi(\mathbf{z}_{t-1}, \xi'_i). \end{aligned}$$

We have that:

$$\begin{aligned} \|\mathbf{z}_{t+1} - \mathbf{z}'_{t+1}\| &\leq \|\tilde{\mathbf{z}}_{t+1} - \tilde{\mathbf{z}}'_{t+1}\| + \|\mathbf{z}_{t+1} - \mathbf{z}'_{t+1} - \tilde{\mathbf{z}}_{t+1} - \tilde{\mathbf{z}}'_{t+1}\| \\ &\leq \|\tilde{\mathbf{z}}_{t+1} - \tilde{\mathbf{z}}'_{t+1}\| + \left\| \frac{\eta}{n} \sum_{i=1}^n (\Xi(\tilde{\mathbf{z}}_{t+1}, \xi_i) - \Xi(\tilde{\mathbf{z}}'_{t+1}, \xi_i) + \Xi(\mathbf{z}'_t, \xi_i) - \Xi(\mathbf{z}_t, \xi_i)) \right\| \\ &\leq (1 + \eta L) \|\tilde{\mathbf{z}}_{t+1} - \tilde{\mathbf{z}}'_{t+1}\| + \eta L (\|\mathbf{z}_t - \mathbf{z}'_t\|). \end{aligned}$$

The growth lemma (restated in Lemma B.1) thus implies:

$$\|\mathbf{z}_{t+1} - \mathbf{z}'_{t+1}\| \leq (1 + \eta L)\xi \|\mathbf{z}_t - \mathbf{z}'_t\| + \left\| \tilde{P}_t(\mathbf{z}'_t) \right\| + \left\| \tilde{P}'_t(\mathbf{z}'_t) \right\| + \eta L (\|\mathbf{z}_t - \mathbf{z}'_t\|),$$

where ξ is the contractiveness of $P + \tilde{P}_t$. Since we are in a strongly monotone and smooth setting, Lemma 2.9 bounds the contractiveness of gradient descent updates by $\xi \leq c_{\text{EG}}$. Because gradient descent has a bounded learning rate and we assume our operator is bounded, we have $\left\| \tilde{P}_t(\mathbf{z}'_t) \right\| \leq \eta K/n$ and $\left\| \tilde{P}'_t(\mathbf{z}'_t) \right\| \leq \eta K/n$. Thus:

$$\|\mathbf{z}_{t+1} - \mathbf{z}'_{t+1}\| \leq ((1 + \eta L)c_{\text{EG}} + \eta L) \|\mathbf{z}_t - \mathbf{z}'_t\| + \frac{2\eta K}{n}.$$

Recurring through $\mathbf{z}_1, \dots, \mathbf{z}_T$:

$$\|\mathbf{z}_T - \mathbf{z}'_T\| \leq \frac{2\eta K}{n} \sum_{t=1}^T ((1 + \eta L)c_{\text{EG}} + \eta L)^{T-t} \leq \frac{2\eta K}{n(1 - ((1 + \eta L)c_{\text{EG}} + \eta L))}$$

where the second inequality is because $\sum_{t=1}^T x^{T-t} = \frac{1-x^T}{1-x}$. \square

B.3 Proof of Theorem 4.2

We restate the relevant result of Klochkov and Zhivotovskiy (2021) below for completeness.

Theorem B.3 (Theorem 1.1 of Klochkov and Zhivotovskiy (2021)). *Let $f: \mathcal{X}^n \times \mathcal{X} \rightarrow [0, M]$ be a data-dependent function that is γ -uniformly stable in its first argument, and that satisfies the Bernstein condition (Assumption 4.4) with parameter B . Suppose the optimization error of f , $\mathbb{E}_{x \sim \mathcal{D}} [f(\mathbf{x}, x)]$, is at most Δ_{opt} . Then there is a universal constant such that, for any $\tilde{c} > 0$ and with probability at least $1 - \delta$ over the randomness of $\mathcal{X} \sim \mathcal{D}^n$,*

$$\mathbb{E}_{x' \sim \mathcal{D}} [f(\mathcal{X}, x')] - \inf_{z^* \in \mathcal{Z}} \mathbb{E}_{x' \sim \mathcal{D}} [f(z^*, x')] \leq \Delta_{\text{opt}} + \tilde{c} \mathbb{E}[\Delta_{\text{opt}}] + c(1 + \frac{1}{\tilde{c}})(\gamma \log n + \frac{M + B}{n}) \log(\frac{1}{\delta}).$$

We can now prove Theorem 4.2, which follows directly from the Bernstein condition bound of Fact 4.5 and Theorem B.3.

Theorem 4.2. *Let (F, \mathcal{Z}) be a smooth strongly monotone variational inequality (Assumptions 2.1, 2.3, 2.5) corresponding to a k -player game. The expected generalization gap for a γ -stable algorithm $A: \mathcal{X}^n \rightarrow \mathcal{Z}$ when training error is small; i.e., $\text{WeakGap}_{\widehat{F}}(A(X)) \leq \varepsilon$, is, with probability at least $1 - \delta$,*

$$\text{WeakGap}_F(A(X)) \leq 2\varepsilon + O\left(\log(\frac{1}{\delta})\left((DL + K \sum_{i \in [k]} \frac{L_i}{\mu_i})\gamma \log n + \frac{(DL + K \sum_{i \in [k]} \frac{L_i}{\mu_i})^2}{n}\right)\right).$$

Proof. We can first apply Theorem B.3 to the data-dependent function $f: \mathcal{X}^n \times \mathcal{X} \rightarrow [0, KD]$ where

$$f(X, x') = \left\langle \widehat{F}_{x'}(A(X)), A(X) - w^*(A(X)) \right\rangle$$

and w^* is as defined in (4). Recall that by summing (8) with (5), we have that f is stable in its first argument:

$$\left| f(X, x') - f(X^{(j)}, x') \right| \leq (2DL + K)\gamma + K\gamma \sum_{i \in [k]} \frac{L_i}{\mu_i}$$

for all $x' \in \mathcal{X}$. Because f also fulfills the Bernstein condition (Fact 4.5), Theorem B.3 implies the main claim. \square