

Training Generative Adversarial Networks via Stochastic Nash Games

Original

Training Generative Adversarial Networks via Stochastic Nash Games / Franci, B., Grammatico, S.. - In: IEEE TRANSACTIONS ON NEURAL NETWORKS AND LEARNING SYSTEMS. - ISSN 2162-237X. - 34:3(2023), pp. 1319-1328. [10.1109/tnnls.2021.3105227]

Availability:

This version is available at: 11583/3003591 since: 2025-10-10T07:49:51Z

Publisher:

Institute of Electrical and Electronics Engineers - IEEE

Published

DOI:10.1109/tnnls.2021.3105227

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

IEEE postprint/Author's Accepted Manuscript

©2023 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collecting works, for resale or lists, or reuse of any copyrighted component of this work in other works.

(Article begins on next page)

Training Generative Adversarial Networks via stochastic Nash games

Barbara Franci¹ and Sergio Grammatico¹

Abstract—Generative adversarial networks (GANs) are a class of generative models with two antagonistic neural networks: a generator and a discriminator. These two neural networks compete against each other through an adversarial process that can be modeled as a stochastic Nash equilibrium problem. Since the associated training process is challenging, it is fundamental to design reliable algorithms to compute an equilibrium.

In this paper, we propose a stochastic relaxed forward-backward (SRFB) algorithm for GANs and we show convergence to an exact solution when an increasing number of data is available. We also show convergence of an averaged variant of the SRFB algorithm to a neighborhood of the solution when only few samples are available. In both cases, convergence is guaranteed when the pseudogradient mapping of the game is monotone. This assumption is among the weakest known in the literature. Moreover, we apply our algorithm to the image generation problem.

Index Terms—Generative Adversarial Networks, stochastic Nash equilibrium problems, variational inequalities, two-player game.

I. INTRODUCTION

A. Generative Adversarial Networks

Generative adversarial networks (GANs) are an example of unsupervised generative model. The basic idea is that, given some samples drawn from a probability distribution, the neural network takes a training set and learns how to obtain an estimate of such distribution. Most of the literature on GANs focuses on sample generation (especially image generation), but they can also be designed to explicitly estimate a probability distribution [1], [2], [3].

The learning process of the neural networks in GANs is made via an adversarial process, in which not only the generative model, but also the opponent, are simultaneously trained. Indeed, there are two neural network classes: the generator that creates data according to a given distribution, and the discriminator that tries to recognize if the samples come from the training data or from the generator. As an example, the generator can be considered as a team of counterfeiters, trying to produce fake currency, while the discriminative model, i.e., the police, tries to detect the counterfeit money [2]. To succeed in this game, the former must learn to reproduce money that are indistinguishable from the original currency, while the discriminator must recognize the samples that are drawn from the same distribution as the training data. Through

the competition, both teams improve their methods until the counterfeit currency is indistinguishable from the original.

Besides this simplistic interpretation, the subject has been widely studied in the literature, because it has many and various applications. In addition to the classic image generation problem [1], [4], GANs have been applied in medicine, e.g., to improve the diagnostic performance of the low-dose computed tomography method [5] and recently to detect pneumonia in potential Covid-19 patients [6]. Moreover, they can be used for correcting images taken under adverse weather conditions (as rain) [7], [8], editing facial attributes [9], image inpainting [10], [11] as well as Pacman [12].

B. Stochastic Nash equilibrium problems

The reason why these networks are called *adversarial* is related to the fact that they can be modeled as a game, where each agent payoff depends on the variables of the other agent [13], [14]. However, the players in GANs can be also considered as cooperative players since they share information with each other [1], [14], [15]. Since there are only the generator and the discriminator, the problem is an instance of a two-player game and it can be also casted as a zero-sum game, depending on choice of the cost functions. From a more general point of view, the class of games that suits the GAN problem is that of stochastic Nash equilibrium problems (SNEPs) where each agent tries to minimize its expected value cost function. Given their connection with game theory, GANs have received theoretical attention as well, both on the study of the associated Nash equilibrium problem [14], [16] and on the design of algorithms to improve the learning and training process [16], [17].

Among the available methods to solve a SNEP, an elegant approach is to recast the problem as a stochastic variational inequality (SVI) [17], [18], [19]. The advantage of this approach is that there are many algorithms available for finding a solution of an SVI, some of them already applied to GANs [17], [20]. For instance, the most used in machine learning is the forward-backward algorithm [21], also known as gradient descent [22], which has the disadvantage that, to ensure convergence, the mapping should be cocoercive, i.e., strongly monotone and Lipschitz continuous. Since the GAN mapping is often non-convex [15], [17], one would prefer an algorithm that is guaranteed to converge for at most monotone mappings. In this case, it is possible to consider the extragradient (EG) algorithm [23], [24], [25] and the forward-backward-forward (FBF) algorithm [26]. The main downside of these two methods is that they require two costly evaluations of the

¹The authors are with the Delft Center for System and Control, TU Delft, The Netherlands {b.franci-1, s.grammatico}@tudelft.nl

This work was partially supported by NWO under research projects OMEGA (613.001.702) and P2P-TALES (647.003.003), and by the ERC under research project COSMOS (802348).

pseudogradient mapping, which is computationally expensive. Due to the large-scale problem size, the ideal algorithm should not be computationally demanding and it should be guaranteed to converge under non-restrictive assumption on the pseudogradient mapping.

C. Contribution

Motivated by the need for computationally light algorithms converging under weak assumptions, we propose an algorithm that requests only one computation of the pseudogradient mapping at each iteration and we show its convergence under mere monotonicity. Specifically, our contributions are summarized next.

- We propose a stochastic relaxed forward-backward (SRFB) algorithm and a variant with averaging (aSRFB) for the training process of GANs. The SRFB involves only one evaluation of the pseudogradient mapping at each iteration, therefore it is computationally cheaper than the EG and FBF algorithms.
- We prove its convergence for monotone mappings, which is considered the “weakest possible” assumption on the pseudogradient mapping [27]. Specifically, whenever only a finite number of samples are available, we prove almost sure convergence to a neighborhood of the solution, while if an increasing set of samples is available, then the algorithm reaches an equilibrium almost surely.
- We apply our algorithm to the image generation problem and compare it with the extragradient scheme.

Our SRFB algorithm is inspired by [28], [29] and a preliminary heuristic application to GAN was presented in [30]. Therein, we do not prove convergence of the SRFB algorithm nor of its aSRFB variant. Moreover, in [30], we only run numerical simulations on synthetic toy examples while in this paper we train the two neural networks for the popular image generation problem with real benchmark data.

D. Related work

Due to the connection between SNEPs and SVIs, many algorithms for variational inequalities have been applied to GANs [17], [19].

The first one is the forward-backward (FB) algorithm [31], [21], also known as gradient descent [22]. It is the most used, even if in many cases it has been proven to be non-convergent [17], [32]. From an operator-theoretic perspective, the FB is not convergent because the pseudogradient mapping should be cocoercive [33] and this is almost never the case in GANs. From an algorithmic perspective, the iterates typically cycle in a neighbourhood of a solution without reaching it [32].

Therefore, research has focused on the forward-backward-forward (FBF) algorithm and on the extragradient (EG) algorithm, that are guaranteed to converge for merely monotone mappings. The FBF algorithm, first presented in [34] and extended to the stochastic case in [26], involves two evaluations of the pseudogradient mapping. A first attempt to apply the FBF algorithm for GANs is presented, along with a relaxed inertial FBF algorithm, in [35]. The extragradient (EG) method

was first proposed in [36] and extended many years later to the stochastic case in [23], [24] and to GANs in [17], [25]. The EG algorithm requires two evaluations of the pseudogradient mapping as well, therefore in [17], a variation is proposed. This involves an extrapolation from the past, i.e., it uses the evaluation of the mapping at previous time steps. In [17] the authors propose also the FB and the EG algorithms with averaging.

The averaging technique was first proposed for VIs in [22] and studied more recently in [17], [37]. In [37], the authors examine two different techniques for averaging: the moving average, which computes the time-average of the iterates, and the exponential moving average which computes an exponentially discounted sum. For both the techniques, they show that, despite convergence cannot be proven, the averaging may help stabilizing the iterates, driving them towards a neighborhood of the solution.

While [37] has mostly a heuristic approach, theoretical convergence studies are presented in [38], [32]. Therein, the authors show that local convergence and stability properties of GAN training depends on the eigenvalues of the Jacobian of the associated gradient vector field.

Another theoretical aspect that has not been extensively addressed yet is the inherent relation between GANs and game theory. In [14], the authors formally introduce Generative Adversarial Network Games, describing (and seeking for) the Nash equilibria of the zero-sum game as saddle points in mixed strategies. The study of saddle-point problems is also studied, in connection with GANs, in [20]. The authors in [15], instead, prove that Adam [39], a second order method for GANs, converges to a stationary local Nash equilibrium.

E. Notation

Let \mathbb{R} indicate the set of real numbers and let $\bar{\mathbb{R}} = \mathbb{R} \cup \{+\infty\}$. $\langle \cdot, \cdot \rangle : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$ denotes the standard inner product and $\| \cdot \|$ represents the associated Euclidean norm. Given N vectors $x_1, \dots, x_N \in \mathbb{R}^n$, $\mathbf{x} := \text{col}(x_1, \dots, x_N) = [x_1^\top, \dots, x_N^\top]^\top$. For a closed set $C \subseteq \mathbb{R}^n$, the mapping $\text{proj}_C : \mathbb{R}^n \rightarrow C$ denotes the projection onto C , i.e., $\text{proj}_C(x) = \text{argmin}_{y \in C} \|y - x\|$.

II. GENERATIVE ADVERSARIAL NETWORKS

The idea behind generative adversarial networks (GANs) is to set up an antagonistic training process between the generator and the discriminator. Typically, the generator and the discriminator are represented by two deep neural networks, and accordingly, they are denoted by two functions, differentiable with respect to their inputs and parameters. The generator creates samples that aim at resembling the distribution of the training data. The generator is therefore trained to fool the discriminator who, in turn, examines the samples to determine whether they are real or fake. This adversarial mechanism can be modeled as a game where the generator and the discriminator represent the players, who want to improve their *payoff* [14].

Formally, the generator is a neural network class, represented by a differentiable function g , with parameters vector

$x_g \in \Omega_g \subseteq \mathbb{R}^{n_g}$. Let us denote the (fake) output of the generator with $g(z, x_g) \in \mathbb{R}^q$ where the input z is a random noise vector drawn from the data prior distribution, $z \sim p_z$ [14]. In game-theoretic terms, the *strategies* of the generator are the parameters x_g that allow g to generate the fake output.

Similarly to the generator, the discriminator is a neural network class with parameter vector $x_d \in \Omega_d \subseteq \mathbb{R}^{n_d}$ and a single output $d(v, x_d) \in [0, 1]$ that indicates how good is the input v . The output of the discriminator can be interpreted as the probability of being real that d assigns to an element v . The *strategies* of the discriminator are the parameters x_d .

Usually [17], [2], the payoff of the discriminator is given by the function

$$J_d(x_g, x_d) = \mathbb{E}[\phi(d(\cdot, x_d))] - \mathbb{E}[\phi(d(g(\cdot, x_g), x_d))], \quad (1)$$

where $\phi : [0, 1] \rightarrow \mathbb{R}$ is a measuring function. The typical choices for ϕ are the Kullback-Leibler divergence or the Jensen-Shannon divergence (a logarithm) as in [2] but other options (such as the Wasserstein distance) are proposed in the literature [3], [40]. Regardless, the mapping in (1) can be interpreted as the distance between the fake value and the real one. The payoff of the generator (J_g) instead, depends on how we describe the game. In fact, the problem can be modeled as a two-player game, or as a zero-sum game, depending on the cost functions. To cast the problem as a zero-sum game, the functions J_g and J_d should satisfy the following relation:

$$J_g(x_g, x_d) = -J_d(x_g, x_d). \quad (2)$$

Then, we can rewrite it as a minmax problem, i.e.,

$$\min_{x_g} \max_{x_d} J_d(x_g, x_d). \quad (3)$$

In words, (3) means that the generator aims at minimizing the distance between the real value and the fake one, while the discriminator wants to maximize such a distance, i.e., d aims at recognizing the generated data.

When the generator has a different payoff function from the discriminator, e.g., given by [17]

$$J_g(x_g, x_d) = \mathbb{E}[\phi(d(g(\cdot, x_g), x_d))], \quad (4)$$

then the problem is not a zero-sum game.

Since the two-player game with cost functions (1) and (4) and the zero-sum game with cost function (1) and relation (2) have the same pseudogradient mapping (defined Section III), it can be proven that the two equilibria are strategically equivalent [14, Th. 10].

III. STOCHASTIC NASH EQUILIBRIUM PROBLEMS

In this section, let us describe the GAN game as a generic stochastic Nash equilibrium problem (SNEP).

The two neural network classes are indexed by the set $\mathcal{I} = \{g, d\}$. Each agent $i \in \mathcal{I}$ has a decision variable $x_i \in \Omega_i \subseteq \mathbb{R}^{n_i}$. In general, the local cost function of agent $i \in \mathcal{I}$ is defined as

$$\mathbb{J}_i(x_i, x_j) = \mathbb{E}_\xi[J_i(x_i, x_j, \xi(\omega))], \quad (5)$$

for some measurable function $J_i : \mathbb{R}^n \times \mathbb{R}^d \rightarrow \mathbb{R}$ where $n = n_d + n_g$. The cost function \mathbb{J}_i of agent $i \in \mathcal{I}$ depends

on its local variable x_i , the decisions of the other player x_j , $j \neq i$, and the random variable $\xi : \Xi \rightarrow \mathbb{R}^d$ that represent the uncertainty. The latter arises when we do not know the distribution of the random variable or it is computationally too expensive to compute the expected value. In practice, this means that we have access only to a finite number of samples from the data distribution. Given the probability space $(\Xi, \mathcal{F}, \mathbb{P})$, \mathbb{E}_ξ indicates the mathematical expectation with respect to the distribution of the random variable $\xi(\omega)$ ¹. Let us suppose that $\mathbb{E}[J_i(x, \xi)]$ is well defined for all the feasible $\mathbf{x} = \text{col}(x_g, x_d) \in \Omega = \Omega_g \times \Omega_d$. For our theoretical analysis, some assumptions on the cost function and the feasible set should be postulated. The following assumptions are standard in monotone game theory [41], [42].

Assumption 1. For each $i \in \mathcal{I}$, the set Ω_i is nonempty, compact and convex.

For each $i, j \in \mathcal{I}$, $i \neq j$, the function $\mathbb{J}_i(\cdot, x_j)$ is convex and continuously differentiable. For each $i \in \mathcal{I}$, $j \neq i$ and for each $\xi \in \Xi$, the function $J_i(\cdot, x_j, \xi)$ is convex, continuously differentiable and Lipschitz continuous with the constant $\ell_i(x_j, \xi)$ integrable in ξ . The function $J_i(x_i, x_j, \cdot)$ is measurable for each x_j , $j \neq i$. ■

Given the decision variable of the other agent, the aim of each agent i is to choose a strategy x_i that solves its local optimization problem, i.e.,

$$\forall i \in \mathcal{I} : \min_{x_i \in \Omega_i} \mathbb{J}_i(x_i, x_j). \quad (6)$$

The solution of the coupled optimization problems in (6) that we are seeking is a stochastic Nash equilibrium (SNE) [42].

Definition 1. A stochastic Nash equilibrium is a collective strategy $\mathbf{x}^* = \text{col}(x_g^*, x_d^*) \in \Omega$ such that for all $i \in \mathcal{I}$

$$\mathbb{J}_i(x_i^*, x_j^*) \leq \inf\{\mathbb{J}_i(y, x_j^*) \mid y \in \Omega_i\}.$$

In other words, a SNE is a pair of strategies where neither the generator, nor the discriminator, can decrease its cost function by unilaterally deviating from its decision.

While existence of a SNE of the game in (6) is guaranteed, under Assumption 1 [42, Section 3.1], uniqueness does not hold in general [42, Section 3.2].

To seek for a Nash equilibrium, we rewrite the problem as a stochastic variational inequality (SVI). Let us first denote the pseudogradient mapping as

$$\mathbb{F}(\mathbf{x}) = \begin{bmatrix} \mathbb{E}[\nabla_{x_g} J_g(x_g, x_d)] \\ \mathbb{E}[\nabla_{x_d} J_d(x_d, x_g)] \end{bmatrix}. \quad (7)$$

We note that the possibility to exchange the expected value and the pseudogradient is ensured by Assumption 1 [42].

Then, the associated SVI reads as

$$\langle \mathbb{F}(\mathbf{x}^*), \mathbf{x} - \mathbf{x}^* \rangle \geq 0 \text{ for all } \mathbf{x} \in \Omega. \quad (8)$$

Remark 1. If Assumption 1 holds, then $\mathbf{x}^* \in \Omega$ is a Nash equilibrium of the game in (6) if and only if \mathbf{x}^* is a solution of the SVI in (8) [18, Prop. 1.4.2], [42, Lem. 3.3].

Moreover, under Assumption 1, the solution set of SVI(Ω, \mathbb{F}) is non empty and compact, i.e., $\text{SOL}(\Omega, \mathbb{F}) \neq \emptyset$ [18, Corollary 2.2.5] and an equilibrium exists. ■

¹From now on, simplicity, we use ξ instead of $\xi(\omega)$ and \mathbb{E} instead of \mathbb{E}_ξ .

Algorithm 1: Stochastic Relaxed forward-backward (SRFB)

- 1 Initialization: $x_i^0 \in \Omega_i$
- 2 Iteration k : Agent i receives x_j^k for $j \neq i$, then updates:

$$\bar{x}_i^k = (1 - \delta)x_i^k + \delta\bar{x}_i^{k-1} \quad (11a)$$

$$x_i^{k+1} = \text{proj}_{\Omega_i}[\bar{x}_i^k - \lambda_i F_i^{\text{VR}}(x_i^k, x_j^k, \xi_i^k)] \quad (11b)$$

In light of Remark 1, we call variational equilibria (v-SNE) the solution of the SVI(Ω, \mathbb{F}) in (8) where \mathbb{F} is as in (7), i.e., the solution of the SVI that are also SNE.

IV. STOCHASTIC RELAXED FORWARD-BACKWARD ALGORITHMS

In this section, we propose two algorithms for solving the SNEP associated to the GANs process: a stochastic relaxed forward backward (SRFB) algorithm and its variant with averaging (aSRFB). The iterations read as in Algorithm 1 and Algorithm 2, respectively and they represent the steps for each agent $i \in \{g, d\}$.

Algorithm 1 and Algorithm 2 differ, besides the presence of the averaging step, on the choice of the approximation used for the pseudogradient mapping. Moreover, we note that the averaging step in Algorithm 2, namely,

$$\mathbf{X}^K = \frac{\sum_{k=1}^K \lambda_k \mathbf{x}^k}{S_K}, \quad S_K = \sum_{k=1}^K \lambda_k \quad (9)$$

can be implemented in a first-order fashion as

$$\mathbf{X}_K = (1 - \tilde{\lambda}_K)\mathbf{X}^{K-1} + \tilde{\lambda}_K \mathbf{x}^K \quad (10)$$

for some $\tilde{\lambda}_K \in [0, 1]$. Moreover, let us remark that (10) is different from (11a) and (14a). Indeed, in Algorithms 1 and 2, (11a) and (14a) are convex combinations, with a constant parameter δ , of the two previous iterates \mathbf{x}^k and $\bar{\mathbf{x}}^{k-1}$, while the averaging in (10) is a weighted cumulative sum over the decision variables \mathbf{x}^k for all the iterations $k \in \{1, \dots, K\}$, with time-varying weights $\{\tilde{\lambda}_k\}_{k=1}^K$. The parameter $\tilde{\lambda}_K$ can be tuned to obtain uniform, geometric or exponential averaging [17], [37].

Let us now describe the approximation schemes used in the definitions of the algorithms. In the SVI framework, there are two main possibilities, depending on the samples available.

Using a finite, fixed number of samples is called stochastic approximation (SA) [21] and it is widely used in the literature of SVIs, in conjunction with conditions on the step sizes to control the stochastic error [24], [43]. In fact, unless the step size sequence is diminishing, it is only possible to prove convergence to a neighborhood of a solution. The stochastic approximation of the pseudogradient mapping, given one sample of the random variable reads as

$$F^{\text{SA}}(\mathbf{x}, \xi) := \begin{bmatrix} \nabla_{x_g} J_g(x_g, x_d, \xi_g) \\ \nabla_{x_d} J_d(x_d, x_g, \xi_d) \end{bmatrix}. \quad (12)$$

F^{SA} uses one or a finite number, called mini-batch, of realizations of the random variable.

Algorithm 2: Stochastic Relaxed forward-backward with averaging (aSRFB)

- 1 Initialization: $x_i^0 \in \Omega_i$
- 2 Iteration $k \in \{1, \dots, K\}$: Agent i receives x_j^k for $j \neq i$, then updates:

$$\bar{x}_i^k = (1 - \delta)x_i^k + \delta\bar{x}_i^{k-1} \quad (14a)$$

$$x_i^{k+1} = \text{proj}_{\Omega_i}[\bar{x}_i^k - \lambda_i F_i^{\text{SA}}(x_i^k, x_j^k, \xi_i^k)] \quad (14b)$$

- 3 Iteration K : $X_i^K = \frac{\sum_{k=1}^K \lambda_k x_i^k}{\sum_{k=1}^K \lambda_k}$
-

When a huge number of samples is available, one can consider using a different approximation scheme, i.e.,

$$F^{\text{VR}}(\mathbf{x}, \xi^k) = \begin{bmatrix} \frac{1}{N_k} \sum_{s=1}^{N_k} \nabla_{x_g} J_i(x_g^k, x_d^k, \xi_g^{(s)}) \\ \frac{1}{N_k} \sum_{s=1}^{N_k} \nabla_{x_d} J_i(x_d^k, x_g^k, \xi_d^{(s)}) \end{bmatrix}. \quad (13)$$

In this case, an increasing number of samples, the batch size N_k , is taken at each iteration [23]. The superscript VR stands for variance reduction and it is related to the property of the approximation error discussed in Remark 2.

V. CONVERGENCE ANALYSIS

A. Basic technical assumptions

With the aim of proving convergence to a solution (or to a neighborhood of one) of Algorithms 1 and 2, we start this section with some assumptions that are common to both the algorithms.

The following monotonicity assumption on the pseudogradient mapping is standard for SVI problems [23], [29], also when applied to GANs [17] and it is the weakest possible to hope for global convergence.

Assumption 2. \mathbb{F} in (7) is monotone, i.e., $\langle \mathbb{F}(\mathbf{x}) - \mathbb{F}(\mathbf{y}), \mathbf{x} - \mathbf{y} \rangle \geq 0$ for all $\mathbf{x}, \mathbf{y} \in \Omega$. ■

Let us now define the filtration $\mathcal{F} = \{\mathcal{F}_k\}$, that is, a family of σ -algebras such that $\mathcal{F}_0 = \sigma(X_0)$, $\mathcal{F}_k = \sigma(X_0, \xi_1, \xi_2, \dots, \xi_k)$ for all $k \geq 1$, and $\mathcal{F}_k \subseteq \mathcal{F}_{k+1}$ for all $k \geq 0$. For all $k \geq 0$, let us also define the stochastic error as

$$\epsilon_k = \hat{F}(\mathbf{x}^k, \xi^k) - \mathbb{F}(\mathbf{x}^k), \quad (15)$$

where \hat{F} indicates one of the two possible approximation schemes. In words, ϵ_k in (15) is the distance between the approximation and the exact expected value mapping. Then, let us postulate that the stochastic error has zero mean and bounded variance, as usual in SVI [17], [23], [29].

Assumption 3. The stochastic error is such that, for all $k \geq 0$, a.s., $\mathbb{E}[\epsilon^k | \mathcal{F}_k] = 0$ and $\mathbb{E}[|\epsilon^k|^2 | \mathcal{F}_k] \leq \sigma^2$. ■

B. Convergence of Algorithm 1

We now state the convergence result for Algorithm 1. First, let us postulate some assumptions functional to our analysis. We start with the batch size sequence, which should be increasing to control the stochastic error.

Assumption 4. The batch size sequence $(N_k)_{k \geq 1}$ is such that, for some $b, k_0, a > 0$, $N_k \geq b(k + k_0)^{a+1}$. ■

Remark 2. Given F^{VR} as in (13), it can be proven that, for some $C > 0$,

$$\mathbb{E} [\|\epsilon_k\|^2 | \mathcal{F}_k] \leq \frac{C\sigma^2}{N_k},$$

i.e., the error diminishes as the batch size increases. Such result is, therefore, called variance reduction. More details can be found in [23, Lemma 3.12] [31, Lemma 6]. ■

In addition to Assumption 2, we postulate that the pseudo-gradient mapping is Lipschitz continuous.

Assumption 5. \mathbb{F} as in (7) is ℓ -Lipschitz continuous for $\ell > 0$, i.e., $\|\mathbb{F}(\mathbf{x}) - \mathbb{F}(\mathbf{y})\| \leq \ell \|\mathbf{x} - \mathbf{y}\|$ for all $\mathbf{x}, \mathbf{y} \in \Omega$. ■

Using the variance reduced scheme in (13), we can take a constant step size, as long as it is small enough while the relaxation parameter should not be too small.

Assumption 6. The step size in Algorithm 1 is such that $\lambda \in (0, \frac{1}{2\delta(2\ell+1)}]$ where ℓ is the Lipschitz constant of \mathbb{F} in (7) as in Assumption 5. The relaxation parameter in Algorithm 1 is such that $\delta \in [\frac{2}{1+\sqrt{5}}, 1]$. ■

We can finally state our first convergence result.

Theorem 1. Let Assumptions 1–6 hold. Then, the sequence $(\mathbf{x}^k)_{k \in \mathbb{N}}$ generated by Algorithm 1 with F^{VR} as in (13) converges a.s. to a SNE of the game in (6).

Proof. See Appendix B. □

C. Convergence of Algorithm 2

In this section, we state the convergence result (and the required assumptions) for Algorithm 2.

First, the bound on the relaxation parameter is wider in this case (compared to Assumption 6).

Assumption 7. The relaxation parameter in Algorithm 2 is such that $\delta \in [0, 1]$. ■

Next, we postulate an assumption on the SA approximation in (12), reasonable in our game theoretic framework [17]. We also assume an explicit bound on the feasible set.

Assumption 8. F^{SA} in (12) is bounded, i.e., there exists $B > 0$ such that for $\mathbf{x} \in \Omega$, $\mathbb{E}[\|F^{\text{SA}}(\mathbf{x}, \xi)\|^2 | \mathcal{F}_k] \leq B$. ■

Assumption 9. The local constraint set Ω is such that $\max_{\mathbf{x}, \mathbf{y} \in \Omega} \|\mathbf{x} - \mathbf{y}\|^2 \leq R^2$, for some $R \geq 0$. ■

To measure how close a point is to the solution, let us introduce the gap function,

$$\text{err}(\mathbf{x}) = \max_{\mathbf{x}^* \in \Omega} \langle \mathbb{F}(\mathbf{x}^*), \mathbf{x} - \mathbf{x}^* \rangle, \quad (16)$$

which is equal 0 if and only if \mathbf{x} is a solution of the (S)VI in (8) [18, Eq. 1.5.2]. Other possible measure functions can be found in [17].

We are now ready to state our second convergence result.

Theorem 2. Let Assumptions 1–3 and 7–9 hold. Let $\mathbf{X}^K = \frac{1}{K} \sum_{k=1}^K \mathbf{x}^k$, $c = \frac{2-\delta^2}{1-\delta}$, B be as in Assumption 8, R be as in Assumption 9 and σ^2 be as in Assumption 3. Then, the sequence $(\mathbf{x}^k)_{k \in \mathbb{N}}$ generated by Algorithm 2 with constant step size and F^{SA} as in (12) is such that

$$\mathbb{E}[\text{err}(\mathbf{X}^K)] = \frac{cR}{\lambda K} + (2B^2 + \sigma^2)\lambda.$$

Thus, $\lim_{K \rightarrow \infty} \mathbb{E}[\text{err}(\mathbf{X}^K)] = (2B^2 + \sigma^2)\lambda$.

Proof. See Appendix C. □

Remark 3. The average defined in Theorem 2 is not in conflict with the definition in (9) because if we consider a fixed step size, it holds that

$$\mathbf{X}^K = \frac{\sum_{k=1}^K \lambda_k \mathbf{x}^k}{\sum_{k=1}^K \lambda_k} = \frac{\lambda \sum_{k=1}^K \mathbf{x}^k}{K\lambda} = \frac{1}{K} \sum_{k=1}^K \mathbf{x}^k. \quad \blacksquare$$

VI. NUMERICAL SIMULATIONS

Let us present some numerical experiments to validate the analysis. We show how GANs are trained using our SRFB algorithm and we propose a comparison with one of the most used algorithms for GANs. Specifically, we compare our SRFB algorithm with the extragradient (EG) algorithm (Algorithm 3) [17]. We note that, compared to Algorithm 1, Algorithm 3 involves two projection steps and two evaluations of the pseudogradient mapping. For the simulations we use Adam (Algorithm 4) [39] instead of the stochastic gradient [21]. In Algorithm 5 we propose the Relaxed Adam, i.e., the SRFB algorithm with Adam; the EG algorithm with Adam can be derived similarly [17, Algorithm 4]. All the simulations are performed on Matlab R2020a with 128G RAM and 2 * Intel(R) Xeon(R) Gold 6148 CPU @ 2.40GHz (20 cores each).

We train two DCGAN architectures [3], [44] (presented in Table I) on the CIFAR10 dataset [45] with the GAN objective [1], [2]. We choose the hyperparameters of Adam as $\beta_1 = 0.5$ and $\beta_2 = 0.9$. We compute the inception score [46] to have an objective comparison: the higher the inception score, the better the image generation. In Figure 1, we show how the inception score increases with time; the solid lines represent a tracking average over the previous and following 50 values of the inceptions score, which is averaged over 20 runs. The transparent area indicated the maximum and minimum values obtained in the 20 runs.

We note that the SRFB algorithm is computationally less demanding than the EG algorithm. Specifically, in Figure 1, after 24 hours (86400 seconds), the SRFB has performed approximately 130000 iterations while the EG 90000. The averaged aSRFB shows worse performances (after approximately 110000 iterations), but this is to be expected since we have convergence only to a neighborhood of the solution (Theorem 2). In Figure 2, we show the mean to variance ratio (average Inception Score divided by its variance) at each time instant of the three algorithms. As one can see, from Figure

Algorithm 3: Extragradient (EG)

- 1 Initialization: $x_i^0 \in \Omega_i$
- 2 Iteration k : Agent i
- 3 Receives x_j^k for $j \neq i$, then updates:

$$y_i^k = \text{proj}_{\Omega_i}[x^k - \alpha_k \hat{F}_i(x_i^k, x_j^k, \xi_i^k)]$$

Receives y_j^k for $j \neq i$, then updates:

$$x_i^{k+1} = \text{proj}_{\Omega_i}[x^k - \alpha_k \hat{F}_i(y_i^k, y_j^k, \xi_i^k)]$$

Algorithm 4: Adam

Input : Initial parameters $x^0, \bar{x}^0 \in \Omega$
 Exponential decay rates $\beta_1, \beta_2 \in [0, 1)$
 Step size α
Output: Parameters x^{k+1}

- 1 **Initialize**
- 2 1st moment vector $z^0 = 0$
- 3 2nd moment vector $y^0 = 0$
- 4 Time step $k=0$
- 5 **for** $k = 1, \dots, K$ **do**
- 6 $\hat{g}^k = \hat{F}(x^k, \xi^k)$ # update pseudogradient
- 7 $z^k = \beta_1 z^{k-1} + (1 - \beta_1) \hat{g}^k$ # update 1st moment
- 8 $y^k = \beta_2 y^{k-1} + (1 - \beta_2) (\hat{g}^k)^2$ # update 2nd moment
- 9 $\tilde{z}^k = \frac{z^k}{1 - \beta_1^k}$ # compute 1st moment estimate
- 10 $\tilde{y}^k = \frac{y^k}{1 - \beta_2^k}$ # compute 2nd moment estimate
- 11 $x^{k+1} = x^k - \alpha \frac{\tilde{z}^k}{\sqrt{\tilde{y}^k + \epsilon}}$ # update parameters
- 12 **end**
- 13 **return** x^{K+1}

Algorithm 5: Relaxed Adam

Input : Initial parameters $x^0, \bar{x}^0 \in \Omega$
 Exponential decay rates $\beta_1, \beta_2 \in [0, 1)$
 Step size α
 Relaxing parameter $\delta \in \left[\frac{2}{1+\sqrt{5}}, 1\right]$
Output: Parameters x^{k+1}

- 1 **Initialize**
- 2 1st moment vector $z^0 = 0$
- 3 2nd moment vector $y^0 = 0$
- 4 Time step $k=0$
- 5 **for** $k = 1, \dots, K$ **do**
- 6 $\hat{g}^k = \hat{F}(x^k, \xi^k)$ # update pseudogradient
- 7 $\bar{x}^k = (1 - \delta)x^k + \delta\bar{x}^{k-1}$ # relaxation step
- 8 $z^k = \beta_1 z^{k-1} + (1 - \beta_1) \hat{g}^k$ # update 1st moment
- 9 $y^k = \beta_2 y^{k-1} + (1 - \beta_2) (\hat{g}^k)^2$ # update 2nd moment
- 10 $\tilde{z}^k = \frac{z^k}{1 - \beta_1^k}$ # compute 1st moment estimate
- 11 $\tilde{y}^k = \frac{y^k}{1 - \beta_2^k}$ # compute 2nd moment estimate
- 12 $x^{k+1} = \bar{x}^k - \alpha \frac{\tilde{z}^k}{\sqrt{\tilde{y}^k + \epsilon}}$ # update parameters
- 13 **end**
- 14 **return** x^K

1 and 2 the SRFB algorithm has a similar performance to the EG algorithm but with a smaller variance, hence a higher reproducibility.

VII. CONCLUSION

The stochastic relaxed forward-backward algorithm is a very promising algorithm for training Generative Adversarial Networks. If an increasing number of samples is available and the pseudogradient mapping of the game is monotone, convergence to the exact solution holds. Instead, with only a finite, fixed mini-batch and the same monotonicity assumption, convergence to a neighborhood of the solution can be proven

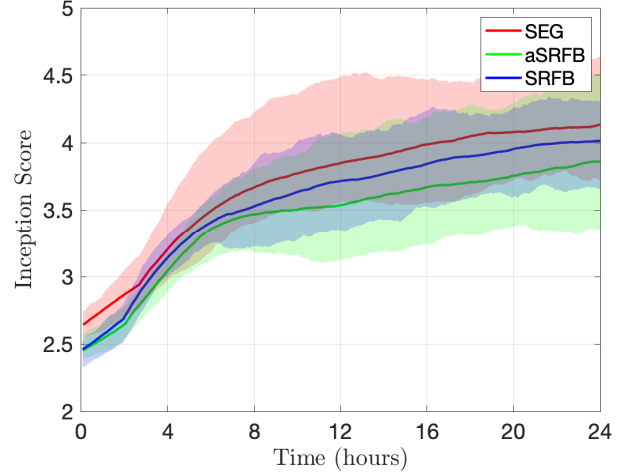


Fig. 1: Inception scores reached by the EG, the SRFB and the aSRFB algorithms.




Fig. 2: Mean to variance ratio corresponding to the average inception scores.

by using an averaging technique. Our numerical experience shows a similar performance compared to the extragradient scheme, widely used in the literature for GANs.

For the future, it would be interesting to extend the convergence result to an exact solution also in the case of a small mini-batch. Since the cost function associated to GAN is often non-convex, it would also worth finding algorithms converging under weaker assumptions than monotonicity.

APPENDIX A PRELIMINARY RESULTS

We here recall some facts about norms, some properties of the projection operator and a preliminary result. Some results find inspiration from [28] where the algorithm is presented in the deterministic case. We start with the norms. We use the cosine rule

$$\langle x, y \rangle = \frac{1}{2}(\langle x, x \rangle + \langle y, y \rangle - \|x - y\|^2) \quad (17)$$

Generator	
Input	$z \in \mathbb{R}^{1 \times 1 \times 100}$
Linear	$100 \rightarrow 5 \times 5 \times 512$
Transp. Conv	(ker: 5×5 , $512 \rightarrow 256$, stride: 1, crop: 0)
	Batch Normalization
	ReLU
Transp. Conv	(ker: 5×5 , $256 \rightarrow 128$, stride: 2, crop: 1)
	Batch Normalization
	ReLU
Transp. Conv	(ker: 5×5 , $128 \rightarrow 3$, stride: 2, crop: 1)
	Tanh(\cdot)
Discriminator	
Input	$z \in \mathbb{R}^{64 \times 64 \times 3}$
Conv	(ker: 5×5 , $3 \rightarrow 64$, stride: 2, pad: 1)
	LeakyReLU
Conv	(ker: 5×5 , $64 \rightarrow 256$, stride: 2, pad: 1)
	Batch Normalization
	LeakyReLU
Conv	(ker: 5×5 , $256 \rightarrow 512$, stride: 2, pad: 1)
	Batch Normalization
	LeakyReLU
Conv	(ker: 5×5 , $512 \rightarrow 1$)

TABLE I: Neural networks used.

and the following two property of the norm [47, Corollary 2.15], $\forall a, b \in \mathcal{E}$, $\forall \alpha \in \mathbb{R}$

$$\|\alpha a + (1 - \alpha)b\|^2 = \alpha\|a\|^2 + (1 - \alpha)\|b\|^2 - \alpha(1 - \alpha)\|a - b\|^2, \quad (18)$$

$$\|a + b\|^2 \leq 2\|a\|^2 + 2\|b\|^2. \quad (19)$$

Concerning the projection operator, by [47, Proposition 12.26], it satisfies the following inequality: let C be a nonempty closed convex set, then, for all $x, y \in C$

$$\bar{x} = \text{proj}_C(x) \Leftrightarrow \langle \bar{x} - x, y - \bar{x} \rangle \geq 0. \quad (20)$$

The projection is also firmly non expansive [47, Prop. 4.16], and consequently, quasi firmly non expansive [47, Def. 4.1].

The Robbins-Siegmund Lemma is widely used in literature to prove a.s. convergence of sequences of random variables.

Lemma 1 (Robbins-Siegmund Lemma, [48]). *Let $\mathcal{F} = (\mathcal{F}_k)_{k \in \mathbb{N}}$ be a filtration. Let $\{\alpha_k\}_{k \in \mathbb{N}}$, $\{\theta_k\}_{k \in \mathbb{N}}$, $\{\eta_k\}_{k \in \mathbb{N}}$ and $\{\chi_k\}_{k \in \mathbb{N}}$ be non negative sequences such that $\sum_k \eta_k < \infty$, $\sum_k \chi_k < \infty$ and let*

$$\forall k \in \mathbb{N}, \quad \mathbb{E}[\alpha_{k+1} | \mathcal{F}_k] + \theta_k \leq (1 + \chi_k)\alpha_k + \eta_k \quad a.s.$$

Then $\sum_k \theta_k < \infty$ and $\{\alpha_k\}_{k \in \mathbb{N}}$ converges a.s. to a non negative random variable. ■

The next lemma collects some properties that follow from the definition of the SRFB algorithm.

Lemma 2. *Given Algorithm 1, the following statements hold.*

- 1) $\mathbf{x}^k - \bar{\mathbf{x}}^{k-1} = \frac{1}{\delta}(\mathbf{x}^k - \bar{\mathbf{x}}^k)$
- 2) $\mathbf{x}^{k+1} - \mathbf{x}^* = \frac{1}{1-\delta}(\bar{\mathbf{x}}^{k+1} - \mathbf{x}^*) - \frac{\delta}{1-\delta}(\bar{\mathbf{x}}^k - \mathbf{x}^*)$
- 3) $\frac{\delta}{(1-\delta)^2} \|\bar{\mathbf{x}}^{k+1} - \bar{\mathbf{x}}^k\|^2 = \delta \|\mathbf{x}^{k+1} - \mathbf{x}^k\|^2$

Proof. Straightforward from Algorithm 1 and [28]. □



(a)



(b)



(c)

Fig. 3: Generated images with the SRFB algorithm (a), with the aSRFB algorithm (b) and with the EG algorithm (c).

APPENDIX B
PROOF OF THEOREM 1

Proof of Theorem 1. Using the property of projection operator (20) we have

$$\langle \mathbf{x}^{k+1} - \bar{\mathbf{x}}^k + \lambda F^{\text{VR}}(\mathbf{x}^k, \xi^k), \mathbf{x}^* - \mathbf{x}^{k+1} \rangle \geq 0, \quad (21)$$

$$\langle \mathbf{x}^k - \bar{\mathbf{x}}^{k-1} + \lambda F^{\text{VR}}(\mathbf{x}^{k-1}, \xi^{k-1}), \mathbf{x}^{k+1} - \mathbf{x}^k \rangle \geq 0. \quad (22)$$

Using Lemma 2.1, (22) becomes

$$\langle \frac{1}{\delta}(\mathbf{x}^k - \bar{\mathbf{x}}^k) + \lambda F^{\text{VR}}(\mathbf{x}^{k-1}, \xi^{k-1}), \mathbf{x}^{k+1} - \mathbf{x}^k \rangle \geq 0. \quad (23)$$

Then, adding (21) and (23) we obtain

$$\begin{aligned} & \langle \mathbf{x}^{k+1} - \bar{\mathbf{x}}^k + \lambda F^{\text{VR}}(\mathbf{x}^k, \xi^k), \mathbf{x}^* - \mathbf{x}^{k+1} \rangle + \\ & + \langle \frac{1}{\delta}(\mathbf{x}^k - \bar{\mathbf{x}}^k) + \lambda F^{\text{VR}}(\mathbf{x}^{k-1}, \xi^{k-1}), \mathbf{x}^{k+1} - \mathbf{x}^k \rangle \geq 0. \end{aligned} \quad (24)$$

Now we use the cosine rule (17),

$$\begin{aligned} & \langle \mathbf{x}^{k+1} - \bar{\mathbf{x}}^k, \mathbf{x}^* - \mathbf{x}^{k+1} \rangle = \\ & -\frac{1}{2}(\|\mathbf{x}^{k+1} - \bar{\mathbf{x}}^k\|^2 + \|\mathbf{x}^{k+1} - \mathbf{x}^*\|^2 - \|\mathbf{x}^* - \bar{\mathbf{x}}^k\|^2) \\ & \langle \frac{1}{\delta}(\mathbf{x}^k - \bar{\mathbf{x}}^k), \mathbf{x}^{k+1} - \mathbf{x}^k \rangle = \\ & -\frac{1}{2\delta}(\|\mathbf{x}^k - \bar{\mathbf{x}}^k\|^2 + \|\mathbf{x}^k - \mathbf{x}^{k+1}\|^2 - \|\mathbf{x}^{k+1} - \bar{\mathbf{x}}^k\|^2) \end{aligned}$$

and we note that

$$\begin{aligned} & \lambda \langle F^{\text{VR}}(\mathbf{x}^k, \xi^k), \mathbf{x}^* - \mathbf{x}^{k+1} \rangle \\ & = -\lambda \langle \mathbb{F}(\mathbf{x}^k), \mathbf{x}^k - \mathbf{x}^* \rangle + \langle \varepsilon^k, \mathbf{x}^* - \mathbf{x}^k \rangle + \\ & + \lambda \langle \mathbb{F}(\mathbf{x}^k), \mathbf{x}^k - \mathbf{x}^{k+1} \rangle + \langle \varepsilon^k, \mathbf{x}^k - \mathbf{x}^{k+1} \rangle. \end{aligned}$$

Then, by reordering and substituting in (24), we obtain:

$$\begin{aligned} & -\|\mathbf{x}^{k+1} - \bar{\mathbf{x}}^k\|^2 - \|\mathbf{x}^{k+1} - \mathbf{x}^*\|^2 + \|\mathbf{x}^* - \bar{\mathbf{x}}^k\|^2 + \\ & -\frac{1}{\delta}\|\mathbf{x}^k - \bar{\mathbf{x}}^k\|^2 - \frac{1}{\delta}\|\mathbf{x}^k - \mathbf{x}^{k+1}\|^2 + \frac{1}{\delta}\|\mathbf{x}^{k+1} - \bar{\mathbf{x}}^k\|^2 + \\ & -2\lambda \langle \mathbb{F}(\mathbf{x}^k), \mathbf{x}^k - \mathbf{x}^* \rangle + 2\lambda \langle \varepsilon^k, \mathbf{x}^* - \mathbf{x}^k \rangle + \\ & + 2\lambda \langle \mathbb{F}(\mathbf{x}^k) - \mathbb{F}(\mathbf{x}^{k-1}), \mathbf{x}^k - \mathbf{x}^{k+1} \rangle + \\ & + 2\lambda \langle \varepsilon^k - \varepsilon^{k-1}, \mathbf{x}^k - \mathbf{x}^{k+1} \rangle \geq 0. \end{aligned} \quad (25)$$

Since \mathbb{F} is monotone, it holds that $\langle \mathbb{F}(\mathbf{x}^k), \mathbf{x}^k - \mathbf{x}^* \rangle \geq \langle \mathbb{F}(\mathbf{x}^*), \mathbf{x}^k - \mathbf{x}^* \rangle \geq 0$. By using Lemma 2.2 and 2.3 as in (28), and substituting in (25), grouping and reordering, we get

$$\begin{aligned} & \frac{1}{1-\delta}\|\bar{\mathbf{x}}^{k+1} - \mathbf{x}^*\|^2 + \frac{1}{\delta}\|\mathbf{x}^k - \mathbf{x}^{k+1}\|^2 \\ & \leq \left(\frac{\delta}{1-\delta} + 1\right)\|\mathbf{x}^* - \bar{\mathbf{x}}^k\|^2 - \frac{1}{\delta}\|\mathbf{x}^k - \bar{\mathbf{x}}^k\|^2 + \\ & + 2\lambda \langle \mathbb{F}(\mathbf{x}^k) - \mathbb{F}(\mathbf{x}^{k-1}), \mathbf{x}^k - \mathbf{x}^{k+1} \rangle \\ & + 2\lambda \langle \varepsilon^k, \mathbf{x}^* - \mathbf{x}^k \rangle + 2\lambda \langle \varepsilon^k - \varepsilon^{k-1}, \mathbf{x}^k - \mathbf{x}^{k+1} \rangle \end{aligned} \quad (26)$$

where we used Assumption 6. Moreover, by using Lipschitz continuity of \mathbb{F} and Cauchy-Schwartz and Young's inequality, it follows that

$$\begin{aligned} & \lambda \langle \mathbb{F}(\mathbf{x}^k) - \mathbb{F}(\mathbf{x}^{k-1}), \mathbf{x}^k - \mathbf{x}^{k+1} \rangle \\ & \leq \frac{\ell\lambda}{2}(\|\mathbf{x}^k - \mathbf{x}^{k-1}\|^2 + \|\mathbf{x}^k - \mathbf{x}^{k+1}\|^2). \end{aligned}$$

Similarly we can bound the term involving the stochastic errors:

$$\begin{aligned} & 2\lambda \langle \varepsilon^k - \varepsilon^{k-1}, \mathbf{x}^k - \mathbf{x}^{k+1} \rangle \\ & \leq 2\lambda \|\varepsilon^k - \varepsilon^{k-1}\| \|\mathbf{x}^k - \mathbf{x}^{k+1}\| \\ & \leq \lambda \|\varepsilon^k - \varepsilon^{k-1}\|^2 + \lambda \|\mathbf{x}^k - \mathbf{x}^{k+1}\|^2. \end{aligned}$$

By substituting in (26), we conclude that

$$\begin{aligned} & \frac{1}{1-\delta}\|\bar{\mathbf{x}}^{k+1} - \mathbf{x}^*\|^2 + \frac{1}{\delta}\|\mathbf{x}^k - \mathbf{x}^{k+1}\|^2 \\ & \leq \frac{1}{1-\delta}\|\mathbf{x}^* - \bar{\mathbf{x}}^k\|^2 - \frac{1}{\delta}\|\mathbf{x}^k - \bar{\mathbf{x}}^k\|^2 + \\ & + \ell\lambda(\|\mathbf{x}^k - \mathbf{x}^{k-1}\|^2 + \|\mathbf{x}^k - \mathbf{x}^{k+1}\|^2) + \\ & + \lambda \|\varepsilon^k - \varepsilon^{k-1}\|^2 + \lambda \|\mathbf{x}^k - \mathbf{x}^{k+1}\|^2 \\ & + 2\lambda \langle \varepsilon^k, \mathbf{x}^* - \mathbf{x}^k \rangle \end{aligned} \quad (27)$$

Now, we consider the residual function of \mathbf{x}^k :

$$\begin{aligned} \text{res}(\mathbf{x}^k)^2 & = \|\mathbf{x}^k - \text{proj}(\mathbf{x}^k - \lambda \mathbb{F}(\mathbf{x}^k))\|^2 \\ & \leq 2\|\mathbf{x}^k - \mathbf{x}^{k+1}\|^2 + 2\|\bar{\mathbf{x}}_k - \mathbf{x}^k + \lambda \varepsilon_k\|^2 \\ & \leq 2\|\mathbf{x}^k - \mathbf{x}^{k+1}\|^2 + 4\|\bar{\mathbf{x}}_k - \mathbf{x}^k\|^2 + 4\lambda^2 \|\varepsilon_k\|^2 \end{aligned}$$

where we added and subtracted $\mathbf{x}^{k+1} = \text{proj}(\bar{\mathbf{x}}_k - \lambda F^{\text{VR}}(\mathbf{x}^k))$ in the first inequality and used the firmly non expansiveness of the projection and (19). It follows that

$$\|\bar{\mathbf{x}}_k - \mathbf{x}^k\|^2 \geq \frac{1}{4} \text{res}(\mathbf{x}^k)^2 - \frac{1}{2} \|\mathbf{x}^k - \mathbf{x}^{k+1}\|^2 - \lambda^2 \|\varepsilon_k\|^2$$

By substituting in (27), we have that

$$\begin{aligned} & \frac{1}{1-\delta}\|\bar{\mathbf{x}}^{k+1} - \mathbf{x}^*\|^2 + \frac{1}{\delta}\|\mathbf{x}^k - \mathbf{x}^{k+1}\|^2 \leq \frac{1}{1-\delta}\|\mathbf{x}^* - \bar{\mathbf{x}}^k\|^2 \\ & - \frac{1}{\delta} \left(\frac{1}{4} \text{res}(\mathbf{x}^k)^2 - \frac{1}{2} \|\mathbf{x}^k - \mathbf{x}^{k+1}\|^2 - \lambda^2 \|\varepsilon_k\|^2 \right) \\ & + \ell\lambda(\|\mathbf{x}^k - \mathbf{x}^{k-1}\|^2 + \|\mathbf{x}^k - \mathbf{x}^{k+1}\|^2) \\ & + \lambda \|\varepsilon^k - \varepsilon^{k-1}\|^2 + \lambda \|\mathbf{x}^k - \mathbf{x}^{k+1}\|^2 + 2\lambda \langle \varepsilon^k, \mathbf{x}^* - \mathbf{x}^k \rangle. \end{aligned}$$

Finally, by taking the expected value, grouping and using Remark 2 and Assumptions 3 and 6, we have

$$\begin{aligned} & \mathbb{E} \left[\frac{1}{1-\delta} \|\bar{\mathbf{x}}^{k+1} - \mathbf{x}^*\|^2 | \mathcal{F}_k \right] + \\ & + \mathbb{E} \left[\left(\frac{1}{2\delta} - \ell\lambda - \lambda \right) \|\mathbf{x}^k - \mathbf{x}^{k+1}\|^2 | \mathcal{F}_k \right] \\ & \leq \frac{1}{1-\delta} \|\mathbf{x}^* - \bar{\mathbf{x}}^k\|^2 + \ell\lambda \|\mathbf{x}^k - \mathbf{x}^{k-1}\|^2 + \\ & + \frac{2\lambda C\sigma}{N_k} + \frac{2\lambda C\sigma}{N_{k-1}} + \frac{\lambda C\sigma}{\delta N_k} \\ & - \frac{1}{\delta} \|\mathbf{x}^k - \bar{\mathbf{x}}^k\|^2 - \frac{1}{4\delta} \text{res}(\mathbf{x}^k)^2. \end{aligned}$$

To use Lemma 1, let $\alpha_k = \frac{1}{1-\delta} \|\mathbf{x}^* - \bar{\mathbf{x}}^k\|^2 + \ell\lambda \|\mathbf{x}^k - \mathbf{x}^{k-1}\|^2$, $\theta_k = \frac{1}{\delta} \|\mathbf{x}^k - \bar{\mathbf{x}}^k\|^2 + \frac{1}{4\delta} \text{res}(\mathbf{x}^k)^2$, $\eta_k = \frac{2\lambda C\sigma}{N_k} + \frac{2\lambda C\sigma}{N_{k-1}} + \frac{\lambda C\sigma}{\delta N_k}$. Applying the Robbins Siegmund Lemma we conclude that α_k converges and that $\sum_{k \in \mathbb{N}} \theta_k$ is summable. This implies that the sequence $(\bar{\mathbf{x}}^k)_{k \in \mathbb{N}}$ is bounded and that $\|\mathbf{x}^k - \bar{\mathbf{x}}^k\| \rightarrow 0$ (otherwise $\sum \frac{1}{\delta} \|\mathbf{x}^k - \bar{\mathbf{x}}^k\|^2 = \infty$). Therefore $(\mathbf{x}^k)_{k \in \mathbb{N}}$ has at least one cluster point $\tilde{\mathbf{x}}$. Moreover, since $\sum_{k \in \mathbb{N}} \theta_k < \infty$, $\text{res}(\mathbf{x}^k)^2 \rightarrow 0$ and $\text{res}(\tilde{\mathbf{x}})^2 = 0$. \square

APPENDIX C
PROOF OF THEOREM 2

Proof of Theorem 2. We start by using the fact that the projection is firmly quasinonexpansive.

$$\begin{aligned} & \|\mathbf{x}^{k+1} - \mathbf{x}^*\|^2 \\ & \leq \|\mathbf{x}^* - \tilde{\mathbf{x}}^k + \lambda F^{\text{SA}}(\mathbf{x}^k, \xi^k)\|^2 \\ & \quad - \|\bar{\mathbf{x}}^k - \lambda F^{\text{SA}}(\mathbf{x}^k, \xi^k) - \mathbf{x}^{k+1}\|^2 \\ & \leq \|\mathbf{x} - \bar{\mathbf{x}}^k\|^2 - \|\bar{\mathbf{x}}^k - \mathbf{x}^{k+1}\|^2 + 2\lambda_k \langle F^{\text{SA}}(\mathbf{x}^k, \xi^k), \mathbf{x}^* - \bar{\mathbf{x}}^k \rangle \\ & \quad + 2\lambda_k \langle F^{\text{SA}}(\mathbf{x}^k, \xi^k), \bar{\mathbf{x}}^k - \mathbf{x}^{k+1} \rangle \\ & = \|\mathbf{x}^* - \bar{\mathbf{x}}^k\|^2 - \|\bar{\mathbf{x}}^k - \mathbf{x}^{k+1}\|^2 + 2\lambda_k \langle F^{\text{SA}}(\mathbf{x}^k, \xi^k), \bar{\mathbf{x}}^k - \mathbf{x}^{k+1} \rangle \\ & \quad + 2\lambda_k \langle F^{\text{SA}}(\mathbf{x}^k, \xi^k), \mathbf{x}^* - \mathbf{x}^k \rangle + 2\lambda_k \langle F^{\text{SA}}(\mathbf{x}^k, \xi^k), \mathbf{x}^k - \bar{\mathbf{x}}^k \rangle \end{aligned}$$

Now we apply Lemma 2.2 and Lemma 2.3 to $\|\mathbf{x}^{k+1} - \mathbf{x}^*\|$:

$$\begin{aligned} \|\mathbf{x}^{k+1} - \mathbf{x}^*\|^2 &= \frac{1}{1-\delta} \|\bar{\mathbf{x}}^{k+1} - \mathbf{x}^*\|^2 - \frac{\delta}{1-\delta} \|\bar{\mathbf{x}}^k - \mathbf{x}^*\|^2 + \\ &\quad + \delta \|\mathbf{x}^{k+1} - \bar{\mathbf{x}}^k\|^2. \end{aligned} \quad (28)$$

Then, we can rewrite the inequality as

$$\begin{aligned} \frac{1}{1-\delta} \|\bar{\mathbf{x}}^{k+1} - \mathbf{x}^*\|^2 &\leq \frac{1}{1-\delta} \|\bar{\mathbf{x}}^k - \mathbf{x}^*\|^2 \\ &+ 2\lambda_k \langle F^{\text{SA}}(\mathbf{x}^k, \xi^k), \mathbf{x}^* - \mathbf{x}^k \rangle + 2\lambda_k \langle F^{\text{SA}}(\mathbf{x}^k, \xi^k), \mathbf{x}^k - \bar{\mathbf{x}}^k \rangle \\ &+ 2\lambda_k \langle F^{\text{SA}}(\mathbf{x}^k, \xi^k), \bar{\mathbf{x}}^k - \mathbf{x}^{k+1} \rangle - (\delta + 1) \|\mathbf{x}^{k+1} - \bar{\mathbf{x}}^k\|^2 \end{aligned} \quad (29)$$

By applying Young's inequality we obtain

$$\begin{aligned} 2\lambda_k \langle F^{\text{SA}}(\mathbf{x}^k, \xi^k), \mathbf{x}^k - \bar{\mathbf{x}}^k \rangle &\leq \lambda_k^2 \|F^{\text{SA}}(\mathbf{x}^k, \xi^k)\|^2 + \|\mathbf{x}^k - \bar{\mathbf{x}}^k\|^2 \\ 2\lambda_k \langle F^{\text{SA}}(\mathbf{x}^k, \xi^k), \bar{\mathbf{x}}^k - \mathbf{x}^{k+1} \rangle &\leq \lambda_k^2 \|F^{\text{SA}}(\mathbf{x}^k, \xi^k)\|^2 + \|\bar{\mathbf{x}}^k - \mathbf{x}^{k+1}\|^2 \end{aligned}$$

Then, inequality (29) becomes

$$\begin{aligned} \frac{1}{1-\delta} \|\bar{\mathbf{x}}^{k+1} - \mathbf{x}^*\|^2 &\leq \frac{1}{1-\delta} \|\bar{\mathbf{x}}^k - \mathbf{x}^*\|^2 \\ &+ 2\lambda_k \langle F^{\text{SA}}(\mathbf{x}^k, \xi^k), \mathbf{x}^* - \mathbf{x}^k \rangle \\ &+ 2\lambda_k^2 \|F^{\text{SA}}(\mathbf{x}^k, \xi^k)\|^2 - (\delta + 1) \|\mathbf{x}^{k+1} - \bar{\mathbf{x}}^k\|^2 \\ &+ \|\mathbf{x}^k - \bar{\mathbf{x}}^k\|^2 + \|\bar{\mathbf{x}}^k - \mathbf{x}^{k+1}\|^2 \end{aligned} \quad (30)$$

Reordering, adding and subtracting $2\lambda_k \langle \mathbb{F}(\mathbf{x}^k), \mathbf{x}^k - \mathbf{x}^* \rangle$ and using Lemma 2, we obtain

$$\begin{aligned} \frac{1}{1-\delta} \|\bar{\mathbf{x}}^{k+1} - \mathbf{x}^*\|^2 + \delta \|\mathbf{x}^{k+1} - \bar{\mathbf{x}}^k\|^2 &\leq \frac{1}{1-\delta} \|\bar{\mathbf{x}}^k - \mathbf{x}^*\|^2 \\ &+ 2\lambda_k \langle \mathbb{F}(\mathbf{x}^k) - F^{\text{SA}}(\mathbf{x}^k, \xi^k), \mathbf{x}^k - \mathbf{x}^* \rangle - 2\lambda_k \langle \mathbb{F}(\mathbf{x}^k), \mathbf{x}^k - \mathbf{x}^* \rangle \\ &+ 2\lambda_k^2 \|F^{\text{SA}}(\mathbf{x}^k, \xi^k)\|^2 + \delta^2 \|\mathbf{x}^k - \mathbf{x}^{k-1}\|^2 \end{aligned} \quad (31)$$

Then, by the definition of ϵ^k , reordering leads to

$$\begin{aligned} 2\lambda_k \langle \mathbb{F}(\mathbf{x}^k), \mathbf{x}^k - \mathbf{x}^* \rangle &\leq \frac{1}{1-\delta} (\|\bar{\mathbf{x}}^k - \mathbf{x}^*\|^2 - \|\bar{\mathbf{x}}^{k+1} - \mathbf{x}^*\|^2) + \\ &+ \delta (\|\mathbf{x}^k - \mathbf{x}^{k-1}\|^2 - \|\mathbf{x}^{k+1} - \bar{\mathbf{x}}^k\|^2) \\ &+ 2\lambda_k^2 \|F^{\text{SA}}(\mathbf{x}^k, \xi^k)\|^2 + 2\lambda_k \langle \epsilon^k, \mathbf{x}^k - \mathbf{x}^* \rangle \end{aligned} \quad (32)$$

Next, we sum over all the iterations, hence inequality (32) becomes

$$\begin{aligned} 2 \sum_{k=1}^K \lambda_k \langle \mathbb{F}(\mathbf{x}^k), \mathbf{x}^k - \mathbf{x}^* \rangle &\leq 2 \sum_{k=1}^K \lambda_k \langle \epsilon^k, \mathbf{x}^k - \mathbf{x}^* \rangle \\ &\leq \frac{1}{1-\delta} \sum_{k=1}^K (\|\bar{\mathbf{x}}^k - \mathbf{x}^*\|^2 - \|\bar{\mathbf{x}}^{k+1} - \mathbf{x}^*\|^2) + \\ &+ \delta \sum_{k=1}^K (\|\mathbf{x}^k - \mathbf{x}^{k-1}\|^2 - \|\mathbf{x}^{k+1} - \bar{\mathbf{x}}^k\|^2) + \\ &+ 2 \sum_{k=1}^K \lambda_k^2 \|F^{\text{SA}}(\mathbf{x}^k, \xi^k)\|^2 \end{aligned} \quad (33)$$

Using Assumption 2 and resolving the sums, we obtain

$$\begin{aligned} 2 \sum_{k=1}^K \lambda_k \langle \mathbb{F}(\mathbf{x}^*), \mathbf{x}^k - \mathbf{x}^* \rangle &\leq 2 \sum_{k=1}^K \lambda_k \langle \epsilon^k, \mathbf{x}^k - \mathbf{x}^* \rangle \\ &\leq \frac{1}{1-\delta} \|\bar{\mathbf{x}}^0 - \mathbf{x}^*\|^2 + \delta \|\mathbf{x}^0 - \bar{\mathbf{x}}^{-1}\|^2 + \\ &+ 2 \sum_{k=1}^K \lambda_k^2 \|F^{\text{SA}}(\mathbf{x}^k, \xi^k)\|^2 \end{aligned} \quad (34)$$

Now, we note that $\langle \epsilon^k, \mathbf{x}^k - \mathbf{x}^* \rangle = \langle \epsilon^k, \mathbf{x}^k - \mathbf{u}^k \rangle + \langle \epsilon^k, \mathbf{u}^k - \mathbf{x}^* \rangle$. We define $\mathbf{u}^0 = \mathbf{x}^0$ and $\mathbf{u}^{k+1} = \text{proj}(\mathbf{u}^k - \lambda_k \epsilon^k)$, thus

$$\begin{aligned} \|\mathbf{u}^{k+1} - \mathbf{x}^*\|^2 &= \|\text{proj}(\mathbf{u}^k - \lambda_k \epsilon^k) - \mathbf{x}^*\|^2 \\ &\leq \|\mathbf{u}^k - \lambda_k \epsilon^k - \mathbf{x}^*\|^2 \\ &\leq \|\mathbf{u}^k - \mathbf{x}^*\|^2 + \lambda_k \|\epsilon^k\|^2 - 2\lambda_k \langle \epsilon^k, \mathbf{u}^k - \mathbf{x}^* \rangle \end{aligned} \quad (35)$$

Therefore, $2\lambda_k \langle \epsilon^k, \mathbf{x}^k - \mathbf{x}^* \rangle = 2\lambda_k \langle \epsilon^k, \mathbf{x}^k - \mathbf{u}^k \rangle + \|\mathbf{u}^k - \mathbf{x}^*\|^2 + \lambda_k \|\epsilon^k\|^2 - \|\mathbf{u}^{k+1} - \mathbf{x}^*\|^2$. By including this in (34) and by doing the sum, we obtain

$$\begin{aligned} 2 \sum_{k=1}^K \lambda_k \langle \mathbb{F}(\mathbf{x}^*), \mathbf{x}^k - \mathbf{x}^* \rangle &\leq \frac{1}{1-\delta} \|\bar{\mathbf{x}}^0 - \mathbf{x}^*\|^2 + \delta \|\mathbf{x}^0 - \bar{\mathbf{x}}^{-1}\|^2 + \\ &+ 2 \sum_{k=1}^K \lambda_k^2 \|F^{\text{SA}}(\mathbf{x}^k, \xi^k)\|^2 + \sum_{k=1}^K \lambda_k^2 \|\epsilon^k\|^2 \\ &+ \|\mathbf{u}_0 - \mathbf{x}^*\|^2 + 2 \sum_{k=1}^K \lambda_k \langle \epsilon^k, \mathbf{x}^k - \mathbf{u}^k \rangle \end{aligned} \quad (36)$$

By definition, $\|\mathbf{u}^0 - \mathbf{x}^*\|^2 = \|\mathbf{x}^0 - \mathbf{x}^*\|^2$. Then, by taking the expected value in (36) and using Assumption 3, we conclude that

$$\begin{aligned} 2 \sum_{k=1}^K \lambda_k \langle \mathbb{F}(\mathbf{x}^*), \mathbf{x}^k - \mathbf{x}^* \rangle &\leq \left(\frac{1}{1-\delta} + 1 \right) \|\bar{\mathbf{x}}^0 - \mathbf{x}^*\|^2 + \delta \|\mathbf{x}^0 - \bar{\mathbf{x}}^{-1}\|^2 + \\ &+ 2 \sum_{k=1}^K \lambda_k^2 \mathbb{E} [\|F^{\text{SA}}(\mathbf{x}^k, \xi^k)\|^2 | \mathcal{F}_k] + \\ &+ \sum_{k=1}^K \lambda_k^2 \mathbb{E} [\|\epsilon^k\|^2 | \mathcal{F}_k]. \end{aligned} \quad (37)$$

Let us define $S = \sum_{k=1}^K \lambda_k$, $\mathbf{X}^K = \frac{\sum_{k=1}^K \lambda_k \mathbf{x}^k}{\sum_{k=1}^K \lambda_k} = \frac{1}{S} \sum_{k=1}^K \lambda_k \mathbf{x}^k$. Then,

$$\begin{aligned} 2S \langle \mathbb{F}(\mathbf{x}^*), \mathbf{X}^K - \mathbf{x}^* \rangle &\leq \frac{2-\delta}{1-\delta} \|\bar{\mathbf{x}}^0 - \mathbf{x}^*\|^2 + \delta \|\mathbf{x}^0 - \bar{\mathbf{x}}^{-1}\|^2 + \\ &+ 2 \sum_{k=1}^K \lambda_k^2 \mathbb{E} [\|F^{\text{SA}}(\mathbf{x}^k, \xi^k)\|^2 | \mathcal{F}_k] + \\ &+ \sum_{k=1}^K \lambda_k^2 \mathbb{E} [\|\epsilon^k\|^2 | \mathcal{F}_k] \\ &\leq \frac{2-\delta^2}{1-\delta} R + (2B^2 + \sigma^2) \sum_{k=1}^K \lambda_k. \end{aligned} \quad (38)$$

Finally, it holds that if λ_k is constant, $S = K\lambda$ and $\sum_{k=1}^K \lambda_k^2 = K\lambda^2$

$$\langle \mathbb{F}(\mathbf{x}^*), \mathbf{X}^K - \mathbf{x}^* \rangle \leq \frac{cR}{K\lambda} + (2B^2 + \sigma^2)\lambda. \quad \square$$

REFERENCES

- [1] I. Goodfellow, "Nips 2016 tutorial: Generative adversarial networks," *arXiv:1701.00160*, 2016.
- [2] I. Goodfellow *et al.*, "Generative adversarial nets," in *Advances in Neural Inf. Processing Syst.*, 2014, pp. 2672–2680.
- [3] Z. Wang, Q. She, and T. E. Ward, "Generative adversarial networks in computer vision: A survey and taxonomy," *arXiv:1906.01529*, 2019.
- [4] X. Song *et al.*, "Sp-gan: Self-growing and pruning generative adversarial networks," *IEEE Trans. Neural Netw. Learn. Syst.*, 2020.
- [5] Q. Yang *et al.*, "Low-dose ct image denoising using a generative adversarial network with wasserstein distance and perceptual loss," *IEEE Trans. Med. Imag.*, vol. 37, no. 6, pp. 1348–1357, 2018.
- [6] N. E. M. Khalifa *et al.*, "Detection of coronavirus (covid-19) associated pneumonia based on generative adversarial networks and a fine-tuned deep transfer learning model using chest x-ray dataset," *arXiv:2004.01184*, 2020.

- [7] H. Zhang, V. Sindagi, and V. M. Patel, "Image de-raining using a conditional generative adversarial network," *IEEE Trans. Circuits Syst. Video Technol.*, 2019.
- [8] P. Xiang *et al.*, "Single-image de-raining with feature-supervised generative adversarial network," *IEEE Signal Process. Lett.*, vol. 26, no. 5, pp. 650–654, 2019.
- [9] K. Zhang *et al.*, "Mu-gan: Facial attribute editing based on multi-attention mechanism," *IEEE/CAA Journal of Automatica Sinica*, 2020.
- [10] Y. Chen *et al.*, "Research on image inpainting algorithm of improved gan based on two-discriminations networks," *Appl. Intell.*, pp. 1–15, 2020.
- [11] —, "The improved image inpainting algorithm via encoder and similarity constraint," *The Vis. Comput.*, pp. 1–15, 2020.
- [12] S. W. Kim *et al.*, "Learning to simulate dynamic environments with gameGAN," in *Proc. IEEE/CVF Conf. Comput. Vis. and Pattern Recognit.*, 2020, pp. 1231–1240.
- [13] S. R. Bulò *et al.*, "Randomized prediction games for adversarial machine learning," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 28, no. 11, pp. 2466–2478, 2016.
- [14] F. A. Oliehoek *et al.*, "GANs: Generative adversarial network games," *arXiv:1712.00679*, 2017.
- [15] M. Heusel *et al.*, "GANs trained by a two time-scale update rule converge to a local nash equilibrium," in *Advances in Neural Inf. Processing Syst.*, 2017, pp. 6626–6637.
- [16] E. Mazumdar, L. J. Ratliff, and S. S. Sastry, "On gradient-based learning in continuous games," *SIAM J. Math. of Data Science*, vol. 2, no. 1, pp. 103–131, 2020.
- [17] G. Gidel *et al.*, "A variational inequality perspective on generative adversarial networks," *arXiv:1802.10551*, 2018.
- [18] F. Facchinei and J.-S. Pang, *Finite-dimensional variational inequalities and complementarity problems*. Springer Sci. & Bus. Media, 2007.
- [19] Q. Tao *et al.*, "Stochastic learning via optimizing the variational inequalities," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 25, no. 10, pp. 1769–1778, 2014.
- [20] P. Mertikopoulos *et al.*, "Optimistic mirror descent in saddle-point problems: Going the extra (gradient) mile," *arXiv:1807.02629*, 2018.
- [21] H. Robbins and S. Monro, "A stochastic approximation method," *Ann. Math. Stat.*, pp. 400–407, 1951.
- [22] R. E. Bruck Jr, "On the weak convergence of an ergodic iteration for the solution of variational inequalities for monotone operators in hilbert space," *J. Math. Anal. Appl.*, vol. 61, no. 1, pp. 159–164, 1977.
- [23] A. Iusem *et al.*, "Extragradient method with variance reduction for stochastic variational inequalities," *SIAM J. Optim.*, vol. 27, no. 2, pp. 686–724, 2017.
- [24] F. Yousefian, A. Nedić, and U. V. Shanbhag, "Optimal robust smoothing extragradient algorithms for stochastic variational inequality problems," in *53rd IEEE Conf. Decis. and Control*. IEEE, 2014, pp. 5831–5836.
- [25] K. Mishchenko *et al.*, "Revisiting stochastic extragradient," in *Int. Conf. Artif. Intell. Statist.* PMLR, 2020, pp. 4573–4582.
- [26] R. I. Boş *et al.*, "Minibatch forward-backward-forward methods for solving stochastic variational inequalities," *Stochastic Syst.*, 2021.
- [27] F. Facchinei and C. Kanzow, "Generalized Nash equilibrium problems," *Ann. Oper. Res.*, vol. 175, no. 1, pp. 177–211, 2010.
- [28] Y. Malitsky, "Golden ratio algorithms for variational inequalities," *Math. Prog.*, Jul 2019. [Online]. Available: <https://doi.org/10.1007/s10107-019-01416-w>
- [29] B. Franci and S. Grammatico, "Stochastic generalized Nash equilibrium seeking in merely monotone games," *arXiv:2002.08318*, 2020.
- [30] —, "A game-theoretic approach for generative adversarial networks," in *59th IEEE Conf. Decis. and Control*. IEEE, 2020, pp. 1646–1651.
- [31] —, "A distributed forward-backward algorithm for stochastic generalized nash equilibrium seeking," *IEEE Trans. Autom. Control*, 2020.
- [32] L. Mescheder, A. Geiger, and S. Nowozin, "Which training methods for GANs do actually converge?" *arXiv:1801.04406*, 2018.
- [33] S. Grammatico, "Comments on "distributed robust adaptive equilibrium computation for generalized convex games"[automatica 63 (2016) 82–91]," *Automatica*, vol. 97, pp. 186–188, 2018.
- [34] P. Tseng, "A modified forward-backward splitting method for maximal monotone mappings," *SIAM J. Control Optim.*, vol. 38, no. 2, pp. 431–446, 2000.
- [35] R. I. Bot, M. Sedlmayer, and P. T. Vuong, "A relaxed inertial forward-backward-forward algorithm for solving monotone inclusions with application to GANs," *arXiv:2003.07886*, 2020.
- [36] G. Korpelevich, "The extragradient method for finding saddle points and other problems," *Matecon*, vol. 12, pp. 747–756, 1976.
- [37] Y. Yazici *et al.*, "The unusual effectiveness of averaging in GAN training," *arXiv:1806.04498*, 2018.
- [38] L. Mescheder, S. Nowozin, and A. Geiger, "The numerics of GANs," in *Advances in Neural Inf. Processing Syst.*, 2017, pp. 1825–1835.
- [39] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv:1412.6980*, 2014.
- [40] M. Arjovsky, S. Chintala, and L. Bottou, "Wasserstein GAN," *arXiv:1701.07875*, 2017.
- [41] F. Facchinei, A. Fischer, and V. Piccialli, "On generalized Nash games and variational inequalities," *Oper. Res. Lett.*, vol. 35, no. 2, pp. 159–164, 2007.
- [42] U. Ravat and U. V. Shanbhag, "On the characterization of solution sets of smooth and nonsmooth convex stochastic Nash games," *SIAM J. Optim.*, 2019.
- [43] J. Koshal, A. Nedic, and U. V. Shanbhag, "Regularized iterative stochastic approximation methods for stochastic variational inequality problems," *IEEE Trans. Autom. Control*, vol. 58, no. 3, pp. 594–609, 2013.
- [44] A. Radford, L. Metz, and S. Chintala, "Unsupervised representation learning with deep convolutional generative adversarial networks," *arXiv:1511.06434*, 2015.
- [45] A. Krizhevsky, G. Hinton *et al.*, "Learning multiple layers of features from tiny images," 2009.
- [46] T. Salimans *et al.*, "Improved techniques for training GANs," in *Advances in Neural Inf. Processing Syst.*, 2016, pp. 2234–2242.
- [47] H. H. Bauschke, P. L. Combettes *et al.*, *Convex analysis and monotone operator theory in Hilbert spaces*. Springer, 2011, vol. 408.
- [48] H. Robbins and D. Siegmund, "A convergence theorem for non negative almost supermartingales and some applications," in *Optim. Methods in Statistics*. Elsevier, 1971, pp. 233–257.



Barbara Franci is a PostDoc at the Delft Center for Systems and Control, Delft University of Technology, Delft, The Netherlands. Her current research interests are on Game Theory and its applications.

She received the Bachelor's and Master's degree in Pure Mathematics from University of Siena, Siena, Italy, respectively in 2012 and 2014. Then, she received her PhD from Politecnico di Torino and University of Turin, Turin, Italy, in 2018. In September - December 2016 she visited the Department of Mechanical Engineering, University of California, Santa Barbara, USA. She was awarded in 2017 with the PhD Quality Award by the Academic Board of Politecnico di Torino.



Sergio Grammatico is an Associate Professor at the Delft Center for Systems and Control, Delft University of Technology, The Netherlands.

Born in Italy, in 1987, he received the Bachelor's degree (summa cum laude) in Computer Engineering, the Master's degree (summa cum laude) in Automatic Control Engineering, and the Ph.D. degree in Automatic Control, all from the University of Pisa, Italy, in February 2008, October 2009, and March 2013 respectively. He also received a Master's degree (summa cum laude) in Engineering

Science from the Sant'Anna School of Advanced Studies, Pisa, Italy, in November 2011. In February-April 2010 and in November-December 2011, he visited the Department of Mathematics, University of Hawaii at Manoa, USA; in January-July 2012, he visited the Department of Electrical and Computer Engineering, University of California at Santa Barbara, USA.

In 2013-2015, he was a post-doctoral Research Fellow in the Automatic Control Laboratory, ETH Zurich, Switzerland. In 2015-2018, he was an Assistant Professor, first in the Department of Electrical Engineering, Control Systems, TU Eindhoven, The Netherlands, then at the Delft Center for Systems and Control, TU Delft, The Netherlands.

He was awarded a 2005 F. Severi B.Sc. Scholarship by the Italian High-Mathematics National Institute, and a 2008 M.Sc. Fellowship by the Sant'Anna School of Advanced Studies. He was awarded 2013 and 2014 "TAC Outstanding Reviewer" by the Editorial Board of the IEEE Trans. on Automatic Control, IEEE Control Systems Society. He was recipient of the Best Paper Award at the 2016 ISDG International Conference on Network Games, Control and Optimization.