
Reducing Noise in GAN Training with Variance Reduced Extragradient

Tatjana Chavdarova*

Mila, Université de Montréal
Idiap, École Polytechnique Fédérale de Lausanne

Gauthier Gidel*

Mila, Université de Montréal
Element AI

François Fleuret

Idiap, École Polytechnique Fédérale de Lausanne

Simon Lacoste-Julien†

Mila, Université de Montréal

Abstract

We study the effect of the stochastic gradient noise on the training of generative adversarial networks (GANs) and show that it can prevent the convergence of standard game optimization methods, while the batch version converges. We address this issue with a novel stochastic variance-reduced extragradient (SVRE) optimization algorithm, which for a large class of games improves upon the previous convergence rates proposed in the literature. We observe empirically that SVRE performs similarly to a batch method on MNIST while being computationally cheaper, and that SVRE yields more stable GAN training on standard datasets.

1 Introduction

Many empirical risk minimization algorithms rely on gradient-based optimization methods. These iterative methods handle large-scale training datasets by computing gradient estimates on a subset of it, a *mini-batch*, instead of using all the samples at each step, the *full batch*, resulting in a method called *stochastic gradient descent* (SGD, [Robbins and Monro \(1951\)](#); [Bottou \(2010\)](#)).

SGD methods are known to efficiently minimize *single* objective loss functions, such as cross-entropy for classification or squared loss for regression. Some algorithms go beyond such training objective and define multiple agents with different or competing objectives. The associated optimization paradigm requires a multi-objective joint minimization. An example of such a class of algorithms are the generative adversarial networks (GANs, [Goodfellow et al., 2014](#)), which aim at finding a Nash equilibrium of a two-player *minimax* game, where the players are deep neural networks (DNNs).

As of their success on supervised tasks, SGD based algorithms have been adopted for GAN training as well. Recently, [Gidel et al. \(2019a\)](#) proposed to use an optimization technique coming from the variational inequality literature called *extragradient* ([Korpelevich, 1976](#)) with provable convergence guarantees to optimize games (see § 2). However, convergence failures, poor performance (sometimes referred to as “mode collapse”), or hyperparameter susceptibility are more commonly reported compared to classical supervised DNN optimization.

We question naive adoption of such methods for game optimization so as to address the reported training instabilities. We argue that as of the two player setting, noise impedes drastically more the training compared to single objective one. More precisely, we point out that the noise due to the stochasticity may break the convergence of the extragradient method, by considering a simplistic stochastic bilinear game for which it provably does *not* converge.

*equal contribution

†Canada CIFAR AI Chair

Method	Complexity	μ -adaptivity	Algorithm 1 Pseudocode for SVRE.
SVRG	$\ln(\frac{1}{\epsilon}) \times (n + \frac{L^2}{\mu^2})$	no	1: Input: Stopping time T , learning rates $\eta_\theta, \eta_\varphi$, initial weights θ_0, φ_0 . $t = 0$
Acc. SVRG	$\ln(\frac{1}{\epsilon}) \times (n + \sqrt{n} \frac{\bar{L}}{\mu})$	no	2: while $t \leq T$ do
SVRE §3.2	$\ln(\frac{1}{\epsilon}) \times (n + \frac{\bar{\ell}}{\mu})$ if $\bar{\ell} = O(\bar{L})$		3: $\varphi^S = \varphi_t$ and $\mu_\varphi^S = \frac{1}{n} \sum_{i=1}^n \nabla_{\varphi} \mathcal{L}_i^D(\theta^S, \varphi^S)$
			4: $\theta^S = \theta_t$ and $\mu_\theta^S = \frac{1}{n} \sum_{i=1}^n \nabla_{\theta} \mathcal{L}_i^G(\theta^S, \varphi^S)$
			5: $N \sim \text{Geom}(1/n)$ (<i>Sample epoch length</i>)
			6: for $i = 0$ to $N-1$ do { <i>Beginning of the epoch</i> }
			7: Sample $i_\theta, i_\varphi \sim \pi_\theta, \pi_\varphi$, do extrapolation:
			8: $\tilde{\varphi}_t = \varphi_t - \eta_\varphi d_{i_\varphi}^D(\theta_t, \varphi_t, \theta^S, \varphi^S) \triangleright (5)$
			9: $\tilde{\theta}_t = \theta_t - \eta_\theta d_{i_\theta}^G(\theta_t, \varphi_t, \theta^S, \varphi^S) \triangleright (5)$
			10: Sample $i_\theta, i_\varphi \sim \pi_\theta, \pi_\varphi$ and do update:
			11: $\varphi_{t+1} = \varphi_t - \eta_\varphi d_{i_\varphi}^D(\tilde{\theta}_t, \tilde{\varphi}_t, \theta^S, \varphi^S) \triangleright (5)$
			12: $\theta_{t+1} = \theta_t - \eta_\theta d_{i_\theta}^G(\tilde{\theta}_t, \tilde{\varphi}_t, \theta^S, \varphi^S) \triangleright (5)$
			13: $t \leftarrow t + 1$
			14: Output: θ_T, φ_T

Table 1: Comparison of variance reduced methods for games for a μ -strongly monotone operator with L_i -Lipschitz stochastic operators. Our result makes the assumption that the operators are ℓ_i -cocoercive. Note that $\ell_i \in [L_i, L_i^2/\mu]$, more details and a tighter rate are provided in §3.2. The SVRG variants are proposed by Palaniappan and Bach (2016). μ -adaptivity indicates if the hyper-parameters that guarantee convergence (step size & epoch length) depend on the strong monotonicity parameter μ : if not, the algorithm is adaptive to local strong monotonicity. Note that in some cases the constant ℓ may depend on μ but SVRE is adaptive to strong convexity when $\bar{\ell}$ remains close to \bar{L} (see for instance Proposition 2).

The theoretical aspect we present in this paper is further supported empirically, since using larger mini-batch sizes for GAN training has been shown to considerably improve the quality of the samples produced by the resulting generative model: Brock et al. (2019) report a relative improvement of 46% of the Inception Score metric (see § 4) on ImageNet if the batch size is increased 8-fold. This notable improvement raises the question if noise reduction optimization methods can be extended to game settings. In turn, this would allow for a principled training method with the practical benefit of omitting to empirically establish this multiplicative factor for the batch size.

In this paper, we investigate the interplay between noise and multi-objective problems in the context of GAN training. Our contributions can be summarized as follows: (i) we show in a motivating example how the noise can make stochastic extragradient fail (see § 2.2). (ii) we propose a new method “stochastic variance reduced extragradient” (SVRE) that combines variance reduction and extrapolation (see Alg. 1 and § 3.2) and show experimentally that it effectively reduces the noise. (iii) we prove the convergence of SVRE under local strong convexity assumptions, improving over the known rates of competitive methods for a large class of games (see § 3.2 for our convergence result and Table 1 for comparison with standard methods). (iv) we test SVRE empirically to train GANs on several standard datasets, and observe that it can improve SOTA deep models in the late stage of their optimization (see § 4).

2 GANs as a Game and Noise in Games

2.1 Game theory formulation of GANs

The models in a GAN are a generator G , that maps an embedding space to the signal space, and should eventually map a fixed noise distribution to the training data distribution, and a discriminator D whose purpose is to allow the training of the generator by classifying genuine samples against generated ones. At each iteration of the algorithm, the discriminator D is updated to improve its “real vs. generated” classification performance, and the generator G to degrade it.

From a game theory point of view, GAN training is a differentiable two-player game where the generator G_θ and the discriminator D_φ aim at minimizing their own cost function \mathcal{L}^G and \mathcal{L}^D , resp.:

$$\theta^* \in \arg \min_{\theta \in \Theta} \mathcal{L}^G(\theta, \varphi^*) \quad \text{and} \quad \varphi^* \in \arg \min_{\varphi \in \Phi} \mathcal{L}^D(\theta^*, \varphi). \quad (2P-G)$$

When $\mathcal{L}^D = -\mathcal{L}^G =: \mathcal{L}$ this game is called a *zero-sum game* and (2P-G) is a minimax problem:

$$\min_{\theta \in \Theta} \max_{\varphi \in \Phi} \mathcal{L}(\theta, \varphi) \quad (SP)$$




Figure 1: Illustration of the discrepancy between games and minimization on simple examples:

$$\min_{\theta, \phi \in \mathbb{R}} \theta^2 + \phi^2, \quad \text{game: } \min_{\theta \in \mathbb{R}} \max_{\phi \in \mathbb{R}} \theta \cdot \phi.$$

Left: Minimization. Up to a neighborhood, the noisy gradient always points to a direction that make the iterate closer to the minimum (\star).
Right: Game. The noisy gradient may point to a direction (red arrow) that push the iterate away from the Nash Equilibrium (\star).

The gradient method does not converge for some convex-concave examples (Mescheder et al., 2017; Gidel et al., 2019a). To address this, Korpelevich (1976) proposed to use the *extragradient* method³ which performs a lookahead step in order to get signal from an *extrapolated* point:

$$\text{Extrapolation: } \begin{cases} \tilde{\theta}_t = \theta_t - \eta \nabla_{\theta} \mathcal{L}^G(\theta_t, \varphi_t) \\ \tilde{\varphi}_t = \varphi_t - \eta \nabla_{\varphi} \mathcal{L}^D(\theta_t, \varphi_t) \end{cases} \quad \text{Update: } \begin{cases} \theta_{t+1} = \theta_t - \eta \nabla_{\theta} \mathcal{L}^G(\tilde{\theta}_t, \tilde{\varphi}_t) \\ \varphi_{t+1} = \varphi_t - \eta \nabla_{\varphi} \mathcal{L}^D(\tilde{\theta}_t, \tilde{\varphi}_t) \end{cases} \quad (\text{EG})$$

Note how θ_t and φ_t are updated with a gradient from a different point, the *extrapolated* one. In the context of a zero-sum game, for any *convex-concave* function \mathcal{L} and any closed convex sets Θ and Φ , the extragradient method converges (Harker and Pang, 1990, Thm. 12.1.11).

2.2 Stochasticity Breaks Extragradient

As the (EG) converges for some examples for which gradient methods do not, it is reasonable to expect that so does its stochastic counterpart (at least to a neighborhood). However, the resulting noise in the gradient estimate may interact in a problematic way with the oscillations due to the *adversarial component* of the game⁴. We depict this phenomenon in Fig. 1, where we show the direction of the noisy gradient on single objective minimization example and contrast it with a multi-objective one.

We present a simplistic example where the extragradient method *converges linearly* (Gidel et al., 2019a, Corollary 1) using the full gradient but *diverges geometrically* when using stochastic estimates of it. Note that standard gradient methods, both batch and stochastic, diverge on this example.

In particular, we show that: (i) if we use standard stochastic estimates of the gradients of \mathcal{L} with a simple finite sum formulation, then the iterates $\omega_t := (\theta_t, \varphi_t)$ produced by the stochastic extragradient method (SEG) diverge geometrically, and on the other hand (ii) the full-batch extragradient method does converge to the Nash equilibrium ω^* of this game (Harker and Pang, 1990, Thm. 12.1.11).

Theorem 1 (Noise may induce divergence). *For any $\epsilon \geq 0$ There exists a zero-sum $\frac{\epsilon}{2}$ -strongly monotone stochastic game such that if $\omega_0 \neq \omega^*$, then for any step-size $\eta > \epsilon$, the iterates (ω_t) computed by the stochastic extragradient method diverge geometrically, i.e., there exists $\rho > 0$, such that $\mathbb{E}[\|\omega_t - \omega^*\|^2] > \|\omega_0 - \omega^*\|^2(1 + \rho)^t$.*

Proof sketch. All detailed proofs can be found in § C of the appendix. We consider the following stochastic optimization problem (with $d = n$):

$$\frac{1}{n} \sum_{i=1}^n \frac{\epsilon}{2} \theta_i^2 + \theta^\top A_i \varphi - \frac{\epsilon}{2} \varphi_i^2 \quad \text{where} \quad [A_i]_{kl} = 1 \text{ if } k = l = i \text{ and } 0 \text{ otherwise.} \quad (1)$$

Note that this problem is a simple dot product between θ and φ with an (ϵ/n) - ℓ_2 norm penalization, thus we can compute the batch gradient and notice that the Nash equilibrium of this problem is $(\theta^*, \varphi^*) = (0, 0)$. However, as we shall see, this simple problem *breaks* with standard stochastic optimization methods.

³For simplicity, we focus on *unconstrained* setting where $\Theta = \mathbb{R}^d$. For the *constrained* case, a Euclidean projection on the constraints set should be added at every update of the method.

⁴Gidel et al. (2019b) formalize the notion of “adversarial component” of a game, which yields a rotational dynamics in gradients methods (oscillations in parameters), as illustrated by the gradient field of Fig. 1 (right).

Sampling a mini-batch without replacement $I \subset \{1, \dots, n\}$, we denote $\mathbf{A}_I := \sum_{i \in I} \mathbf{A}_i$. The extragradient update rule can be written as:

$$\begin{cases} \boldsymbol{\theta}_{t+1} = (1 - \eta \mathbf{A}_I \epsilon) \boldsymbol{\theta}_t - \eta \mathbf{A}_I ((1 - \eta \mathbf{A}_J \epsilon) \boldsymbol{\varphi}_t + \eta \mathbf{A}_J \boldsymbol{\theta}_t) \\ \boldsymbol{\varphi}_{t+1} = (1 - \eta \mathbf{A}_I \epsilon) \boldsymbol{\varphi}_t + \eta \mathbf{A}_I ((1 - \eta \mathbf{A}_J \epsilon) \boldsymbol{\theta}_t - \eta \mathbf{A}_J \boldsymbol{\varphi}_t), \end{cases} \quad (2)$$

where I and J are the mini-batches sampled for the update and the extrapolation step, respectively. Let us write $N_t := \|\boldsymbol{\theta}_t\|^2 + \|\boldsymbol{\varphi}_t\|^2$. Noticing that $[\mathbf{A}_I \boldsymbol{\theta}]_i = [\boldsymbol{\theta}]_i$ if $i \in I$ and 0 otherwise, we have,

$$\mathbb{E}[N_{t+1}] = \left(1 - \frac{|I|}{n} (2\eta\epsilon - \eta^2(1 + \epsilon^2)) - \frac{|I|^2}{n^2} (2\eta^2 - \eta^4(1 + \epsilon^2))\right) \mathbb{E}[N_t]. \quad (3)$$

Consequently, if the mini-batch size is smaller than half of the dataset size, i.e. $2|I| \leq n$, we have that $\forall \eta > \epsilon, \exists \rho > 0, \text{ s.t. }, \mathbb{E}[N_t] > N_0(1 + \rho)^t$. For the theorem statement, we set $n = 2$ and $|I| = 1$.

This result may seem contradictory with the standard result on SEG (Juditsky et al., 2011) saying that the average of the iterates computed by SEG does converge to the Nash equilibrium of the game. However, an important assumption made by Juditsky et al. is that the iterates are projected onto a compact set and that estimator of the gradient has finite variance. These assumptions break in this example since the variance of the estimator is proportional to the norm of the (unbounded) parameters. Note that constraining the optimization problem (23) to bounded domains Θ and Φ , would make the finite variance assumption from Juditsky et al. (2011) holds. Consequently, the averaged iterate $\bar{\omega}_t := \frac{1}{t} \sum_{s=0}^{t-1} \omega_s$ would converge to ω^* . In § A.1, we explain why in a *non-convex setting*, the convergence of the *last iterate* is preferable.

3 Reducing Noise in Games with Variance Reduced Extragradient

One way to reduce the noise in the estimation of the gradient is to use mini-batches of samples instead of one sample. However, mini-batch stochastic extragradient fails to converge on (23) if the mini-batch size is smaller than half of the dataset size (see § C.1). In order to get an estimator of the gradient with a vanishing variance, the optimization literature proposed to take advantage of the finite-sum formulation that often appears in machine learning (Schmidt et al., 2017, and references therein).

3.1 Variance Reduced Gradient Methods

Let us assume that the objective in (2P-G) can be decomposed as a finite sum such that⁵

$$\mathcal{L}^G(\omega) = \frac{1}{n} \sum_{i=1}^n \mathcal{L}_i^G(\omega) \quad \text{and} \quad \mathcal{L}^D(\omega) = \frac{1}{n} \sum_{i=1}^n \mathcal{L}_i^D(\omega) \quad \text{where} \quad \omega := (\boldsymbol{\theta}, \boldsymbol{\varphi}). \quad (4)$$

Johnson and Zhang (2013) propose the “stochastic variance reduced gradient” (SVRG) as an *unbiased* estimator of the gradient with a smaller variance than the vanilla mini-batch estimate. The idea is to occasionally take a snapshot ω^S of the current model’s parameters, and store the full batch gradient μ^S at this point. Computing the full batch gradient μ^S at ω^S is an expensive operation but not prohibitive if done infrequently (for instance once every dataset pass).

Assuming that we have stored ω^S and $\mu^S := (\mu_\theta^S, \mu_\varphi^S)$, the SVRG estimates of the gradients are:

$$\mathbf{d}_i^G(\omega) := \frac{\nabla \mathcal{L}_i^G(\omega) - \nabla \mathcal{L}_i^G(\omega^S)}{n\pi_i} + \mu_\theta^S, \quad \mathbf{d}_i^D(\omega) := \frac{\nabla \mathcal{L}_i^D(\omega) - \nabla \mathcal{L}_i^D(\omega^S)}{n\pi_i} + \mu_\varphi^S. \quad (5)$$

These estimates are unbiased: $\mathbb{E}[\mathbf{d}_i^G(\omega)] = \frac{1}{n} \sum_{i=1}^n \nabla \mathcal{L}_i^G(\omega) = \nabla \mathcal{L}^G(\omega)$, where the expectation is taken over i , picked with probability π_i . The non-uniform sampling probabilities π_i are used to bias the sampling according to the Lipschitz constant of the stochastic gradient in order to sample more often gradients that change quickly. This strategy has been first introduced for variance reduced methods by Xiao and Zhang (2014) for SVRG and has been discussed for saddle point optimization by Palaniappan and Bach (2016).

Originally, SVRG was introduced as an epoch based algorithm with a *fixed epoch size*: in Alg. 1, one epoch is an inner loop of size N (Line 6). However, Hofmann et al. (2015) proposed instead to

⁵The “noise dataset” in a GAN is not finite though; see § D.1 for details on how to cope with this in practice.

sample the size of each epoch from a geometric distribution, enabling them to analyze SVRG the same way as SAGA under a unified framework called q -memorization algorithm. We generalize their framework to handle the extrapolation step (EG) and provide a convergence proof for such q -memorization algorithms for games in § C.2.

One advantage of Hofmann et al. (2015)’s framework is also that the sampling of the epoch size does not depend on the condition number of the problem, whereas the original proof for SVRG had to consider an epoch size larger than the condition number (see Leblond et al. (2018, Corollary 16) for a detailed discussion on the convergence rate for SVRG). Thus, this new version of SVRG with a random epoch size becomes *adaptive to the local strong convexity* since none of its hyper-parameters depend on the strong convexity constant.

However, because of some new technical aspects when working with monotone operators, Palaniappan and Bach (2016)’s proofs (both for SAGA and SVRG) require a step-size (and epoch length for SVRG) that depends on the strong monotonicity constant making these algorithms not adaptive to local strong monotonicity. This motivates the proposed SVRE algorithm, which may be adaptive to local strong monotonicity, and is thus more appropriate for non-convex optimization.

3.2 SVRE: Stochastic Variance Reduced Extragradient

We describe our proposed algorithm called stochastic variance reduced extragradient (SVRE) in Alg. 1. In an analogous manner to how Palaniappan and Bach (2016) combined SVRG with the gradient method, SVRE combines SVRG estimates of the gradient (5) with the *extragradient method* (EG).

With SVRE we are able to improve the convergence rates for variance reduction for a large class of stochastic games (see Table 1 and Thm. 2), and we show in § 3.3 that it is the only method which empirically converges on the simple example of § 2.2.

We now describe the theoretical setup for the convergence result. A standard assumption in convex optimization is the assumption of strong convexity of the function. However, in a game, the operator,

$$v : \omega \mapsto [\nabla_{\theta} \mathcal{L}^G(\omega), \nabla_{\varphi} \mathcal{L}^D(\omega)]^{\top}, \quad (6)$$

associated with the updates is no longer the gradient of a single function. To make an analogous assumption for games the optimization literature considers the notion of *strong monotonicity*.

Definition 1. An operator $F : \omega \mapsto (F_{\theta}(\omega), F_{\varphi}(\omega)) \in \mathbb{R}^{d+p}$ is said to be $(\mu_{\theta}, \mu_{\varphi})$ -strongly monotone if for all $\omega, \omega' \in \mathbb{R}^{p+d}$ we have

$$\Omega((\theta, \varphi), (\theta', \varphi')) := \mu_{\theta} \|\theta - \theta'\|^2 + \mu_{\varphi} \|\varphi - \varphi'\|^2 \leq (F(\omega) - F(\omega'))^{\top} (\omega - \omega'),$$

where we write $\omega := (\theta, \varphi) \in \mathbb{R}^{d+p}$. A monotone operator is a $(0, 0)$ -strongly monotone operator.

This definition is a generalization of strong convexity for operators: if f is μ -strongly convex, then ∇f is a μ -monotone operator. Another assumption is the γ regularity assumption,

Definition 2. An operator $F : \omega \mapsto (F_{\theta}(\omega), F_{\varphi}(\omega)) \in \mathbb{R}^{d+p}$ is said to be $(\gamma_{\theta}, \gamma_{\varphi})$ -regular if,

$$\gamma_{\theta}^2 \|\theta - \theta'\|^2 + \gamma_{\varphi}^2 \|\varphi - \varphi'\|^2 \leq \|F(\omega) - F(\omega')\|^2, \quad \forall \omega, \omega' \in \mathbb{R}^{p+d}. \quad (7)$$

Note that an operator is always $(0, 0)$ -regular. This assumption originally introduced by Tseng (1995) has been recently used (Azizian et al., 2019) to improve the convergence rate of extragradient. For instance for a full rank bilinear matrix problem γ is its smallest singular value. More generally, in the case $\gamma_{\theta} = \gamma_{\varphi}$, the regularity constant is a lower bound on the minimal singular value of the Jacobian of F (Azizian et al., 2019).

One of our main assumptions is the cocoercivity assumption, which implies the Lipschitzness of the operator in the unconstrained case. We use the cocoercivity constant because it provides a tighter bound for general strongly monotone and Lipschitz games (see discussion following Theorem 2).

Definition 3. An operator $F : \omega \mapsto (F_{\theta}(\omega), F_{\varphi}(\omega)) \in \mathbb{R}^{d+p}$ is said to be $(\ell_{\theta}, \ell_{\varphi})$ -cocoercive, if for all $\omega, \omega' \in \Omega$ we have

$$\|F(\omega) - F(\omega')\|^2 \leq \ell_{\theta} (F_{\theta}(\omega) - F_{\theta}(\omega'))^{\top} (\theta - \theta') + \ell_{\varphi} (F_{\varphi}(\omega) - F_{\varphi}(\omega'))^{\top} (\varphi - \varphi'). \quad (8)$$

Note that for a L -Lipschitz and μ -strongly monotone operator, we have $\ell \in [L, L^2/\mu]$ (Facchinei and Pang, 2003). For instance, when F is the gradient of a convex function, we have $\ell = L$. More generally, when $F(\omega) = (\nabla f(\theta) + M\varphi, \nabla g(\varphi) - M^\top \theta)$, where f and g are μ -strongly convex and L smooth we have that $\gamma = \sigma_{\min}(M)$ and $\|M\|^2 = O(\mu L)$ is a sufficient condition for $\ell = O(L)$ (see §B). Under this assumption on each cost function of the game operator, we can define a cocoercivity constant adapted to the non-uniform sampling scheme of our stochastic algorithm:

$$\bar{\ell}(\pi)^2 := \frac{1}{n} \sum_{i=1}^n \frac{1}{n\pi_i} \ell_i^2. \quad (9)$$

The standard *uniform sampling scheme* corresponds to $\pi_i := \frac{1}{n}$ and the optimal *non-uniform sampling scheme* corresponds to $\tilde{\pi}_i := \frac{\ell_i}{\sum_{i=1}^n \ell_i}$. By Jensen's inequality, we have: $\bar{\ell}(\tilde{\pi}) \leq \bar{\ell}(\pi) \leq \max_i \ell_i$.

For our main result, we make strong convexity, cocoercivity and regularity assumptions.

Assumption 1. For $1 \leq i \leq n$, the gradients $\nabla_{\theta} \mathcal{L}_i^G$ and $\nabla_{\varphi} \mathcal{L}_i^D$ are respectively ℓ_i^{θ} and ℓ_i^{φ} -cocoercive and $(\gamma_i^{\theta}, \gamma_i^{\varphi})$ -regular. The operator (6) is $(\mu_{\theta}, \mu_{\varphi})$ -strongly monotone.

We now present our convergence result for SVRE with non-uniform sampling (to make our constants comparable to those of Palaniappan and Bach (2016)), but note that we have used uniform sampling in all our experiments (for simplicity).

Theorem 2. Under Assumption 1, after t iterations, the iterate $\omega_t := (\theta_t, \varphi_t)$ computed by SVRE (Alg. 1) with step-size $\eta_{\theta} \leq (40\bar{\ell}_{\theta})^{-1}$ and $\eta_{\varphi} \leq (40\bar{\ell}_{\varphi})^{-1}$ and sampling scheme $(\tilde{\pi}_{\theta}, \tilde{\pi}_{\varphi})$ verifies:

$$\mathbb{E}[\|\omega_t - \omega^*\|_2^2] \leq \left(1 - \frac{1}{2} \min \left\{ \eta_{\theta} \mu_{\theta} + \frac{9\eta_{\theta}^2 \bar{\gamma}_{\theta}^2}{10}, \eta_{\varphi} \mu_{\varphi} + \frac{9\eta_{\varphi}^2 \bar{\gamma}_{\varphi}^2}{10}, \frac{4}{5n} \right\} \right)^t \mathbb{E}[\|\omega_0 - \omega^*\|_2^2],$$

where $\bar{\ell}_{\theta}(\pi_{\theta})$ and $\bar{\ell}_{\varphi}(\pi_{\varphi})$ are defined in (9). Particularly, for $\eta_{\theta} = \frac{1}{40\bar{\ell}_{\theta}}$ and $\eta_{\varphi} = \frac{1}{40\bar{\ell}_{\varphi}}$ we get

$$\mathbb{E}[\|\omega_t - \omega^*\|_2^2] \leq \left(1 - \frac{1}{2} \min \left\{ \frac{1}{40} \left(\frac{\mu_{\theta}}{\bar{\ell}_{\theta}} + \frac{\bar{\gamma}_{\theta}^2}{45\bar{\ell}_{\theta}^2} \right), \frac{1}{40} \left(\frac{\mu_{\varphi}}{\bar{\ell}_{\varphi}} + \frac{\bar{\gamma}_{\varphi}^2}{45\bar{\ell}_{\varphi}^2} \right), \frac{4}{5n} \right\} \right)^t \mathbb{E}[\|\omega_0 - \omega^*\|_2^2].$$

We prove this theorem in § C.2. We can notice that the respective *condition numbers* of \mathcal{L}^G and \mathcal{L}^D defined as $\kappa_{\theta} := \frac{\mu_{\theta}}{\bar{\ell}_{\theta}} + \frac{\bar{\gamma}_{\theta}^2}{\bar{\ell}_{\theta}^2}$ and $\kappa_{\varphi} := \frac{\mu_{\varphi}}{\bar{\ell}_{\varphi}} + \frac{\bar{\gamma}_{\varphi}^2}{\bar{\ell}_{\varphi}^2}$ appear in our convergence rate. The cocoercivity constant ℓ belongs to $[L, L^2/\mu]$, thus our rate may be significantly faster⁶ than the convergence rate of the (non-accelerated) algorithm of Palaniappan and Bach (2016) that depends on the product $\frac{\mu_{\theta}}{L_{\theta}} \frac{\mu_{\varphi}}{L_{\varphi}}$. They avoid a dependence on the maximum of the condition numbers squared, $\max\{\kappa_{\varphi}^2, \kappa_{\theta}^2\}$, by using the weighted Euclidean norm $\Omega(\theta, \varphi)$ defined in (14) and rescaling the functions \mathcal{L}^G and \mathcal{L}^D with their strong-monotonicity constant. However, this rescaling trick suffers from two issues: (i) we do not know in practice a good estimate of the strong monotonicity constant, which was *not* the case in Palaniappan and Bach (2016)'s application; and (ii) the algorithm does not adapt to local strong-monotonicity. This property is important in non-convex optimization since we want the algorithm to exploit the (potential) local stability properties of a stationary point.

3.3 Motivating example

The example (23) for $\epsilon = 0$ seems to be challenging in the stochastic setting since all the standard methods and even the stochastic extragradient method fails to find its Nash equilibrium (note that this example is *not* strongly monotone). We set $n = d = 100$, and draw $[A_i]_{kl} = \delta_{kli}$ and $[b_i]_k, [c_i]_k \sim \mathcal{N}(0, 1/d)$, $1 \leq k, l \leq d$, where $\delta_{kli} = 1$ if $k = l = i$ and 0 otherwise. Our optimization problem is:

$$\min_{\theta \in \mathbb{R}^d} \max_{\varphi \in \mathbb{R}^d} \frac{1}{n} \sum_{i=1}^n (\theta^\top b_i + \theta^\top A_i \varphi + c_i^\top \varphi). \quad (10)$$

⁶Particularly, when F is the gradient of a convex function (or close to it) we have $\ell \approx L$ and thus our rate recovers the standard $\ln(1/\epsilon)L/\mu$, improving over the accelerated algorithm of Palaniappan and Bach (2016). More generally, under the assumptions of Proposition 2, we also recover $\ln(1/\epsilon)L/\mu$.

We compare variants of the following algorithms (with uniform sampling and average our results over 5 different seeds): (i) AltSGD: the standard method to train GANs—stochastic gradient with alternating updates of each player. (ii) SVRE: Alg. 1. The AVG prefix correspond to the *uniform average* of the iterates, $\bar{\omega} := \frac{1}{t} \sum_{s=0}^{t-1} \omega_s$. We observe in Fig. 4 that AVG-SVRE converges sublinearly (whereas AVG-AltSGD fails to converge).

This motivates a new variant of SVRE based on the idea that even if the averaged iterate converges, we do not compute the gradient at that point and thus we do not benefit from the fact that this iterate is closer to the optimum (see § A.1). Thus the idea is to occasionally restart the algorithm, i.e., consider the averaged iterate as the new starting point of our algorithm and compute the gradient at that point. Restart goes well with SVRE as we already occasionally stop the inner loop to recompute μ^S , at which point we decide (with a probability p to be fixed) whether or not to restart the algorithm by taking the snapshot at point $\bar{\omega}_t$ instead of ω_t . This variant of SVRE is described in Alg. 3 in § E and the variant combining VRad in § D.1.

In Fig. 4 we observe that the only method that converges is SVRE and its variants. We do not provide convergence guarantees for Alg. 3 and leave its analysis for future work. However, it is interesting that, to our knowledge, this algorithm is the only stochastic algorithm (excluding batch extragradient as it is not stochastic) that converge for (23). Note that we tried all the algorithms presented in Fig. 3 from Gidel et al. (2019a) on this *unconstrained* problem and that all of them diverge.

4 GAN Experiments

In this section, we investigate the empirical performance of SVRE for *GAN training*. Note, however, that our theoretical analysis does not hold for games with non-convex objectives such as GANs.

Datasets. We used the following datasets: (i) **MNIST** (Lecun and Cortes), (ii) **CIFAR-10** (Krizhevsky, 2009, §3), (iii) **SVHN** (Netzer et al., 2011), and (iv) **ImageNet ILSVRC 2012** (Russakovsky et al., 2015), using 28×28 , $3 \times 32 \times 32$, $3 \times 32 \times 32$, and $3 \times 64 \times 64$ resolution, respectively.

Metrics. We used the **Inception score** (IS, Salimans et al., 2016) and the **Fréchet Inception distance** (FID, Heusel et al., 2017) as performance metrics for image synthesis. To gain insights if SVRE indeed reduces the variance of the gradient estimates, we used the **second moment estimate–SME** (uncentered variance), computed with an exponentially moving average. See § F.1 for details.

DNN architectures. For experiments on **MNIST**, we used the DCGAN architectures (Radford et al., 2016), described in § F.2.1. For real-world datasets, we used two architectures (see § F.2 for details and § F.2.2 for motivation): (i) **SAGAN** (Zhang et al., 2018), and (ii) **ResNet**, replicating the setup of Miyato et al. (2018), described in detail in § F.2.3 and F.2.4, respectively. For clarity, we refer the former as *shallow*, and the latter as *deep* architectures.

Optimization methods. We conduct experiments using the following optimization methods for GANs: (i) **BatchE**: full-batch extragradient, (ii) **SG**: stochastic gradient (alternating GAN), and (iii) **SE**: stochastic extragradient, and (iv) **SVRE**: stochastic variance reduced extragradient. These can be combined with adaptive learning rate methods such as *Adam* or with parameter averaging, hereafter denoted as **–A** and **AVG–**, respectively. In § D.1, we present a variant of Adam adapted to variance reduced algorithms, that is referred to as **–VRAd**. When using the SE–A baseline and *deep* architectures, the convergence rapidly fails at some point of training (cf. § G.3). This motivates experiments where we start from a stored checkpoint taken *before* the baseline diverged, and *continue training with SVRE*. We denote these experiments with **WS–SVRE** (warm-start SVRE).

4.1 Results

Comparison on MNIST. The **MNIST** common benchmark allowed for comparison with full-batch extragradient, as it is feasible to compute. Fig. 3 depicts the IS metric while using either a stochastic, full-batch or variance reduced version of extragradient (see details of SVRE-GAN in § D.2). We always combine the stochastic baseline (SE) with *Adam*, as proposed by Gidel et al. (2019a). In terms of number of parameter updates, SVRE performs similarly to BatchE–A (see Fig. 5a, § G). Note that the latter requires significantly more computation: Fig. 3a depicts the IS metric using the number of mini-batch computations as x-axis (a surrogate for the wall-clock time, see below). We observe that,

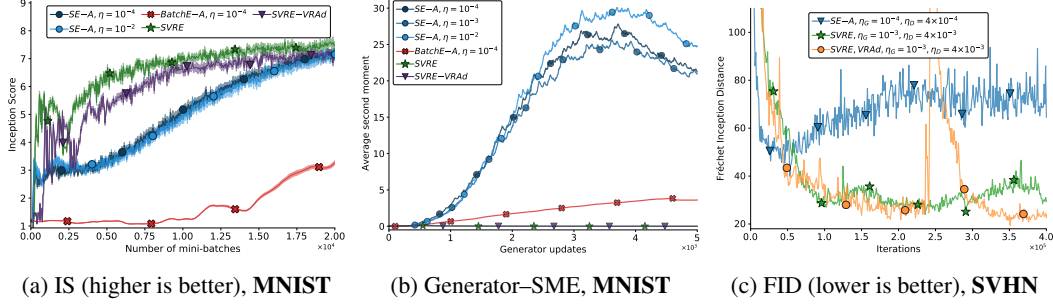


Figure 3: **Figures a & b.** Stochastic, full-batch and variance reduced extragradient optimization on **MNIST**. We used $\eta = 10^{-2}$ for SVRE. SE-A with $\eta = 10^{-3}$ achieves similar IS performances as $\eta = 10^{-2}$ and $\eta = 10^{-4}$, omitted from Fig. a for clarity. **Figure c.** FID on **SVHN**, using *shallow* architectures. See § 4 and § F for naming of methods and details on the implementation, respectively.




Figure 4: Distance to the optimum of (10), see § 3.3 for the experimental setup.

	SG-A	SE-A	SVRE	WS-SVRE
CIFAR-10	21.70	18.65	23.56	16.77
SVHN	5.66	5.14	4.81	4.88

Table 2: Best obtained FID scores for the different optimization methods using the *deep* architectures (see Table 8, § F.2.4). WS-SVRE starts from the best obtained scores of SE-A. See § F and § G for implementation details and additional results, respectively.

as SE-A has slower per-iteration convergence rate, SVRE converges faster on this dataset. At the end of training, all methods reach similar performances (IS is above 8.5, see Table 9, § G).

Computational cost. The relative cost of one pass over the dataset for SVRE versus vanilla SGD is a factor of 5: the full batch gradient is computed (on average) after one pass over the dataset, giving a slowdown of 2; the factor 5 takes into account the extra stochastic gradient computations for the variance reduction, as well as the extrapolation step overhead. However, as SVRE provides less noisy gradient, it may converge faster per iteration, compensating the extra per-update cost. Note that many computations can be done in parallel. In Fig. 3a, the x-axis uses an implementation-independent surrogate for wall-clock time that counts the number of mini-batch gradient computations. Note that some training methods for GANs require multiple discriminator updates per generator update, and we observed that to stabilize our baseline when using the *deep* architectures it was required to use 1:5 update ratio of $G:D$ (cf. § G.3), whereas for SVRE we used ratio of 1:1 (Tab. 2 lists the results).

Second moment estimate and Adam. Fig. 3b depicts the averaged second-moment estimate for parameters of the Generator, where we observe that SVRE effectively reduces it over the iterations. The reduction of these values may be the reason why Adam combined with SVRE performs poorly (as these values appear in the denominator, see § D.1). To our knowledge, SVRE is the first optimization method with a constant step size that has worked empirically for GANs on non-trivial datasets.

Comparison on real-world datasets. In Fig. 3c, we compare SVRE with the SE-A baseline on **SVHN**, using *shallow* architectures. We observe that although SE-A in some experiments obtains better performances in the early iterations, SVRE allows for obtaining improved *final* performances. Tab. 2 summarizes the results on **CIFAR-10** and **SVHN** with *deep* architectures. We observe that, with deeper architectures, SE-A is notably more unstable, as training collapsed in 100% of the experiments. To obtain satisfying results for SE-A, we used various techniques such as a schedule of the learning rate and different update ratios (see § G.3). On the other hand, SVRE *did not collapse in any of the experiments* but took longer time to converge compared to SE-A. Interestingly, although

WS-SVRE starts from an iterate point after which the baseline diverges, it continues to improve the obtained FID score and does not diverge. See § G for additional experiments.

5 Related work

Surprisingly, there exist only a few works on variance reduction methods for monotone operators, namely from [Palaniappan and Bach \(2016\)](#) and [Davis \(2016\)](#). The latter requires a co-coercivity assumption on the operator and thus only convex optimization is considered. Our work provides a new way to use variance reduction for monotone operators, using the extragradient method ([Korpelevich, 1976](#)). Recently, [Iusem et al. \(2017\)](#) proposed an extragradient method with variance reduction for an *infinite sum* of operators. The authors use mini-batches of growing size in order to reduce the variance of their algorithm and to converge with a constant step-size. However, this approach is prohibitively expensive in our application. Moreover, [Iusem et al.](#) are not using the SAGA/SVRG style of updates exploiting the finite sum formulation, leading to sublinear convergence rate, while our method benefits from a linear convergence rate exploiting the finite sum assumption.

[Daskalakis et al. \(2018\)](#) proposed a method called Optimistic-Adam inspired by game theory. This method is closely related to extragradient, with slightly different update scheme. More recently, [Gidel et al. \(2019a\)](#) proposed to use extragradient to train GANs, introducing a method called ExtraAdam. This method outperformed Optimistic-Adam when trained on CIFAR-10. Our work is also an attempt to find principled ways to train GANs. Considering that the game aspect is better handled by the extragradient method, we focus on the optimization issues arising from the noise in the training procedure, a disregarded potential issue in GAN training.

In the context of deep learning, despite some very interesting theoretical results on non-convex minimization ([Reddi et al., 2016](#); [Allen-Zhu and Hazan, 2016](#)), the effectiveness of variance reduced methods is still an open question, and a recent technical report by [Defazio and Bottou \(2018\)](#) provides negative empirical results on the variance reduction aspect. In addition, two recent large scale studies showed that increased batch size has: (i) only marginal impact on *single objective training* ([Shallue et al., 2018](#)) and (ii) a surprisingly large performance improvement on *GAN training* ([Brock et al., 2019](#)). In our work, we are able to show positive results for variance reduction in a real-world deep learning setting. This unexpected difference seems to confirm the remarkable discrepancy, that remains poorly understood, between multi-objective optimization and standard minimization.

6 Discussion

Motivated by a simple bilinear game optimization problem where stochasticity provably breaks the convergence of previous stochastic methods, we proposed the novel SVRE algorithm that combines SVRG with the extragradient method for optimizing games. On the theory side, SVRE improves upon the previous best results for strongly-convex games, whereas empirically, it is the only method that converges for our stochastic bilinear game counter-example.

We empirically observed that SVRE for GAN training obtained convergence speed similar to Batch-Extragradient on MNIST, while the latter is computationally infeasible for large datasets. For shallow architectures, SVRE matched or improved over baselines on all four datasets. Our experiments with deeper architectures show that SVRE is notably more stable with respect to hyperparameter choice. Moreover, while its stochastic counterpart diverged in all our experiments, SVRE did not. However, we observed that SVRE took more iterations to converge when using deeper architectures, though notably, we were using constant step-sizes, unlike the baselines which required Adam. As adaptive step-sizes often provide significant improvements, developing such an appropriate version for SVRE is a promising direction for future work. In the meantime, the stability of SVRE suggests a practical use case for GANs as warm-starting it just before the baseline diverges, and running it for further improvements, as demonstrated with the WS-SVRE method in our experiments.

Acknowledgements

This research was partially supported by the Canada CIFAR AI Chair Program, the Canada Excellence Research Chair in “Data Science for Realtime Decision-making”, by the NSERC Discovery Grant RGPIN-2017-06936, by the Hasler Foundation through the MEMUDE project, and by a Google

Focused Research Award. Authors would like to thank Compute Canada for providing the GPUs used for this research. TC would like to thank Sebastian Stich and Martin Jaggi, and GG and TC would like to thank Hugo Berard for helpful discussions.

References

- Z. Allen-Zhu and E. Hazan. Variance reduction for faster non-convex optimization. In *ICML*, 2016.
- W. Azizian, I. Mitliagkas, S. Lacoste-Julien, and G. Gidel. A tight and unified analysis of extragradient for a whole spectrum of differentiable games. *arXiv preprint arXiv:1906.05945*, 2019.
- L. Bottou. Large-scale machine learning with stochastic gradient descent. In *COMPSTAT*, 2010.
- S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge university press, 2004.
- A. Brock, J. Donahue, and K. Simonyan. Large scale GAN training for high fidelity natural image synthesis. In *ICLR*, 2019.
- C. Daskalakis, A. Ilyas, V. Syrgkanis, and H. Zeng. Training GANs with optimism. In *ICLR*, 2018.
- D. Davis. Smart: The stochastic monotone aggregated root-finding algorithm. *arXiv:1601.00698*, 2016.
- A. Defazio and L. Bottou. On the ineffectiveness of variance reduced optimization for deep learning. *arXiv:1812.04529*, 2018.
- A. Defazio, F. Bach, and S. Lacoste-Julien. Saga: A fast incremental gradient method with support for non-strongly convex composite objectives. In *NIPS*, 2014.
- F. Facchinei and J.-S. Pang. *Finite-Dimensional Variational Inequalities and Complementarity Problems Vol I*. Springer Series in Operations Research and Financial Engineering, Finite-Dimensional Variational Inequalities and Complementarity Problems. Springer-Verlag, 2003.
- G. Gidel, H. Berard, P. Vincent, and S. Lacoste-Julien. A variational inequality perspective on generative adversarial nets. In *ICLR*, 2019a.
- G. Gidel, R. A. Hemmat, M. Pezeshki, R. L. Priol, G. Huang, S. Lacoste-Julien, and I. Mitliagkas. Negative momentum for improved game dynamics. In *AISTATS*, 2019b.
- X. Glorot and Y. Bengio. Understanding the difficulty of training deep feedforward neural networks. In *AISTATS*, 2010.
- I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *NIPS*, 2014.
- P. T. Harker and J.-S. Pang. Finite-dimensional variational inequality and nonlinear complementarity problems: a survey of theory, algorithms and applications. *Mathematical programming*, 1990.
- K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. *arXiv:1512.03385*, 2015.
- M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter. GANs trained by a two time-scale update rule converge to a local nash equilibrium. In *NIPS*, 2017.
- T. Hofmann, A. Lucchi, S. Lacoste-Julien, and B. McWilliams. Variance reduced stochastic gradient descent with neighbors. In *NIPS*, 2015.
- S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, 2015.
- A. Iusem, A. Jofré, R. I. Oliveira, and P. Thompson. Extragradient method with variance reduction for stochastic variational inequalities. *SIAM Journal on Optimization*, 2017.
- R. Johnson and T. Zhang. Accelerating stochastic gradient descent using predictive variance reduction. In *NIPS*, 2013.

- A. Juditsky, A. Nemirovski, and C. Tauvel. Solving variational inequalities with stochastic mirror-prox algorithm. *Stochastic Systems*, 2011.
- D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015.
- G. Korpelevich. The extragradient method for finding saddle points and other problems. *Matecon*, 1976.
- A. Krizhevsky. Learning Multiple Layers of Features from Tiny Images. Master’s thesis, 2009.
- R. Leblond, F. Pedregosa, and S. Lacoste-Julien. Improved asynchronous parallel optimization analysis for stochastic incremental methods. *JMLR*, 19(81):1–68, 2018.
- Y. Lecun and C. Cortes. The MNIST database of handwritten digits. URL <http://yann.lecun.com/exdb/mnist/>.
- J. H. Lim and J. C. Ye. Geometric GAN. *arXiv:1705.02894*, 2017.
- L. Mescheder, S. Nowozin, and A. Geiger. The numerics of GANs. In *NIPS*, 2017.
- T. Miyato, T. Kataoka, M. Koyama, and Y. Yoshida. Spectral normalization for generative adversarial networks. In *ICLR*, 2018.
- Y. Netzer, T. Wang, A. Coates, A. Bissacco, B. Wu, and A. Y. Ng. Reading digits in natural images with unsupervised feature learning. 2011. URL <http://ufldl.stanford.edu/housenumbers/>.
- B. Palaniappan and F. Bach. Stochastic variance reduction methods for saddle-point problems. In *NIPS*, 2016.
- A. Radford, L. Metz, and S. Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. In *ICLR*, 2016.
- S. J. Reddi, A. Hefny, S. Sra, B. Póczos, and A. Smola. Stochastic variance reduction for nonconvex optimization. In *ICML*, 2016.
- H. Robbins and S. Monro. A stochastic approximation method. *The Annals of Mathematical Statistics*, 1951.
- O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *IJCV*, 115(3):211–252, 2015.
- T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen. Improved techniques for training GANs. In *NIPS*, 2016.
- T. Schaul, S. Zhang, and Y. LeCun. No more pesky learning rates. In *ICML*, 2013.
- M. Schmidt, N. Le Roux, and F. Bach. Minimizing finite sums with the stochastic average gradient. *Mathematical Programming*, 2017.
- C. J. Shallue, J. Lee, J. Antognini, J. Sohl-Dickstein, R. Frostig, and G. E. Dahl. Measuring the effects of data parallelism on neural network training. *arXiv:1811.03600*, 2018.
- C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna. Rethinking the inception architecture for computer vision. *arXiv:1512.00567*, 2015.
- P. Tseng. On linear convergence of iterative methods for the variational inequality problem. *Journal of Computational and Applied Mathematics*, 1995.
- A. C. Wilson, R. Roelofs, M. Stern, N. Srebro, and B. Recht. The marginal value of adaptive gradient methods in machine learning. In *NIPS*, 2017.
- L. Xiao and T. Zhang. A proximal stochastic gradient method with progressive variance reduction. *SIAM Journal on Optimization*, 24(4):2057–2075, 2014.
- H. Zhang, I. Goodfellow, D. Metaxas, and A. Odena. Self-Attention Generative Adversarial Networks. *arXiv:1805.08318*, 2018.

A Noise in games

A.1 Why is convergence of the last iterate preferable?

In light of Theorem 1, the behavior of the iterates on the unconstrained version of (23) ($\epsilon = 0$):

$$\min_{\theta \in \Theta} \max_{\varphi \in \Phi} \frac{1}{n} \sum_{i=1}^n \theta^\top A_i \varphi \quad \text{where} \quad [A_i]_{kl} = 1 \text{ if } k = l = i \text{ and } 0 \text{ otherwise.} \quad (11)$$

where Θ and Φ are compact and convex sets, is the following: they will diverge until they reach the boundary of Θ and Φ and then they will start to turn around the Nash equilibrium of (11) lying on these boundaries. Using convexity properties, we can then show that the averaged iterates will converge to the Nash equilibrium of the problem. However, with an arbitrary large domain, this convergence rate may be arbitrary slow (since it depends on the diameter of the domain).

Moreover, this behavior might be even more problematic in a non-convex framework because even if by chance we initialize close to the Nash equilibrium, we would get away from it and we cannot rely on convexity to expect the average of the iterates to converge.

Consequently, we would like optimization algorithms generating iterates that *stay close to the Nash equilibrium*.

B Definitions and Lemmas

B.1 Smoothness and Monotonicity of the operator

Another important property used is the Lipschitzness of an operator.

Definition 4. A mapping $F : \mathbb{R}^p \rightarrow \mathbb{R}^d$ is said to be L -Lipschitz if,

$$\|F(\omega) - F(\omega')\|_2 \leq L \|\omega - \omega'\|_2, \quad \forall \omega, \omega' \in \Omega. \quad (12)$$

Definition 5. A differentiable function $f : \Omega \rightarrow \mathbb{R}$ is said to be μ -strongly convex if

$$f(\omega) \geq f(\omega') + \nabla f(\omega')^\top (\omega - \omega') + \frac{\mu}{2} \|\omega - \omega'\|_2^2 \quad \forall \omega, \omega' \in \Omega. \quad (13)$$

Definition 6. A function $(\theta, \varphi) \mapsto \mathcal{L}(\theta, \varphi)$ is said convex-concave if $\mathcal{L}(\cdot, \varphi)$ is convex for all $\varphi \in \Phi$ and $\mathcal{L}(\theta, \cdot)$ is concave for all $\theta \in \Theta$. An \mathcal{L} is said to be μ -strongly convex concave if $(\theta, \varphi) \mapsto \mathcal{L}(\theta, \varphi) - \frac{\mu}{2} \|\theta\|_2^2 + \frac{\mu}{2} \|\varphi\|_2^2$ is convex concave.

Definition 7. For $\mu_\theta, \mu_\varphi > 0$, an operator $F : \omega \mapsto (F_\theta(\omega), F_\varphi(\omega)) \in \mathbb{R}^{d+p}$ is said to be $(\mu_\theta, \mu_\varphi)$ -strongly monotone if $\forall \omega, \omega' \in \Omega \subset \mathbb{R}^{p+d}$,

$$(F(\omega) - F(\omega'))^\top (\omega - \omega') \geq \mu_\theta \|\theta - \theta'\|^2 + \mu_\varphi \|\varphi - \varphi'\|^2.$$

where we noted $\omega := (\theta, \varphi) \in \mathbb{R}^{d+p}$.

Definition 8. An operator $F : (\omega) \in \mathbb{R}^d$ is said to be ℓ -cocoercive, if for all $\omega, \omega' \in \Omega$ we have

$$\|F(\omega) - F(\omega')\|^2 \leq \ell (F(\omega) - F(\omega'))^\top (\omega - \omega'). \quad (14)$$

Proposition 1 (Folklore). A L -Lipschitz and μ -strongly monotone operator is L^2/μ -cocoercive

Proof. By applying lipschitzness and strong monotonicity,

$$\|F(\omega) - F(\omega')\|^2 \leq L^2 \|\omega - \omega'\|^2 \leq L^2/\mu (F(\omega) - F(\omega'))^\top (\omega - \omega') \quad (15)$$

□

Proposition 2. If $F(\omega) = (\nabla f(\theta) + M\varphi, \nabla g(\varphi) - M^\top \theta)$, where f and g are μ -strongly convex and L smooth, then $\|M\|^2 = O(\mu L)$ is a sufficient condition for F to be ℓ -cocoercive with $\ell = O(L)$

Proof. We rewrite F as the sum of the gradient of convex Lipschitz function F_{grad} and a L -Lipschitz and μ -strongly monotone operator F_{mon} :

$$F_{grad}(\omega) := (\nabla f(\theta) - \mu\theta, \nabla g(\varphi) - \mu\varphi) \quad \text{and} \quad F_{mon} : (M\varphi + \mu\theta, -M^\top \theta + \mu\varphi) \quad (16)$$

Then

$$\|F(\omega) - F(\omega')\|^2 \leq 2\|F_{grad}(\omega) - F_{grad}(\omega')\|^2 + 2\|F_{mon}(\omega) - F_{mon}(\omega')\|^2 \quad (17)$$

$$\leq 2(L + \mu)(F_{grad}(\omega) - F_{grad}(\omega'))^\top (\omega - \omega') \quad (18)$$

$$+ 2(\|M\| + \mu)^2 / \mu (F_{mon}(\omega) - F_{mon}(\omega'))^\top (\omega - \omega') \quad (19)$$

$$= O(L)(F_{grad}(\omega) - F_{grad}(\omega'))^\top (\omega - \omega') \quad (20)$$

$$+ O(L)(F_{mon}(\omega) - F_{mon}(\omega'))^\top (\omega - \omega') \quad (21)$$

$$= O(L)(F(\omega) - F(\omega'))^\top (\omega - \omega') \quad (22)$$

where for the second inequality we used that a $(L + \mu)$ -Lipschitz convex function is $(L + \mu)$ -cocoercive and Proposition 2. \square

C Proof of Theorems

C.1 Proof of Theorem 1

Proof. We consider the following stochastic optimization problem,

$$\frac{1}{n} \sum_{i=1}^n \frac{\epsilon}{2} \theta_i^2 + \theta^\top \mathbf{A}_i \varphi - \frac{\epsilon}{2} \varphi_i^2 = \frac{1}{n} \sum_{i=1}^n \frac{\epsilon}{2} \|\mathbf{A}_i \theta\|^2 + \theta^\top \mathbf{A}_i \varphi - \frac{\epsilon}{2} \|\mathbf{A}_i \varphi\|^2 \quad (23)$$

where $[\mathbf{A}_i]_{kl} = 1$ if $k = l = i$ and 0 otherwise. Note that $(\mathbf{A}_i)^\top = \mathbf{A}_i$ for $1 \leq i \leq n$. Let us consider the extragradient method where to compute an unbiased estimator of the gradients at (θ, φ) we sample $i \in \{1, \dots, n\}$ and use $[\mathbf{A}_i \theta, \mathbf{A}_i \varphi]$ as estimator of the vector flow.

In this proof we note, $\mathbf{A}_I := \sum_{i \in I} \mathbf{A}_i$ and $\theta^{(I)}$ the vector such that $[\theta^{(I)}]_i = [\theta]_i$ if $i \in I$ and 0 otherwise. Note that $\mathbf{A}_I \theta = \theta^{(I)}$ and that $\mathbf{A}_I \mathbf{A}_J = \mathbf{A}_{I \cap J}$.

Thus the extragradient update rule can be noted as

$$\begin{cases} \theta_{t+1} = (1 - \eta \mathbf{A}_I \epsilon) \theta_t - \eta \mathbf{A}_I ((1 - \eta \mathbf{A}_J \epsilon) \varphi_t + \eta \mathbf{A}_J \theta_t) \\ \varphi_{t+1} = (1 - \eta \mathbf{A}_I \epsilon) \varphi_t + \eta \mathbf{A}_I ((1 - \eta \mathbf{A}_J \epsilon) \theta_t - \eta \mathbf{A}_J \varphi_t) \end{cases} \quad (24)$$

where I is the mini-batch sampled (without replacement) for the update and J the mini-batch sampled (without replacement) for the extrapolation.

We can thus notice that, when $I \cap J = \emptyset$, we have

$$\begin{cases} \theta_{t+1} = \theta_t - \eta \epsilon \theta_t^{(I)} - \eta \varphi_t^{(I)} \\ \varphi_{t+1} = \varphi_t - \eta \epsilon \varphi_t^{(I)} + \eta \theta_t^{(I)}, \end{cases} \quad (25)$$

and otherwise,

$$\begin{cases} \theta_{t+1} = \theta_t - \eta \epsilon \theta_t^{(I)} - \eta \varphi_t^{(I)} - \eta^2 (\theta_t^{(I \cap J)} - \epsilon \varphi_t^{(I \cap J)}) \\ \varphi_{t+1} = \varphi_t - \eta \epsilon \varphi_t^{(I)} + \eta \theta_t^{(I)} - \eta^2 (\varphi_t^{(I \cap J)} + \epsilon \theta_t^{(I \cap J)}). \end{cases} \quad (26)$$

The intuition is that, on one hand, when $I \cap J = \emptyset$ (which happens with high probability when $|I| \ll n$, e.g., when $|I| = 1$, $\mathbb{P}(I \cap J = \emptyset) = 1 - 1/n$), the algorithm performs an update that get away from the Nash equilibrium when $2\epsilon \geq \eta$:

$$(25) \Rightarrow N_{t+1} = N_t + (\eta^2 \epsilon^2 + \eta^2 - 2\eta\epsilon) N_t^{(I)}, \quad (27)$$

where $N_t := \|\theta_t\|^2 + \|\varphi_t\|^2$ and $N_t^{(I)} := \|\theta_t^{(I)}\|^2 + \|\varphi_t^{(I)}\|^2$. On the other hand, The updates that provide improvement only happen when $I \cap J$ is large (which happen with low probability, e.g., when $|I| = 1$, $\mathbb{P}(I \cap J \neq \emptyset) = 1/n$):

$$(26) \Rightarrow N_{t+1} = N_t - N_t^{(I)}(2\eta\epsilon - \eta^2(1 + \epsilon^2)) - N_t^{(I \cap J)}(2\eta^2 - \eta^4(1 + \epsilon^2)) \quad (28)$$

Conditioning on θ_t and φ_t , we get that

$$\mathbb{E}[N_t^{(I \cap J)} | \theta_t, \varphi_t] = \sum_{i=1}^n \mathbb{P}(i \in I \cap J) ([\theta_t]_i^2 + [\varphi_t]_i^2) \quad \text{and} \quad \mathbb{P}(i \in I \cap J) = \mathbb{P}(i \in I) \mathbb{P}(i \in J) = \frac{|I|^2}{n^2}. \quad (29)$$

Leading to,

$$\mathbb{E}[N_t^{(I \cap J)} | \theta_t, \varphi_t] = \frac{|I|^2}{n^2} \sum_{i=1}^n ([\theta_t]_i^2 + [\varphi_t]_i^2) = \frac{|I|^2}{n^2} N_t \quad \text{and} \quad \mathbb{E}[N_t^{(I)} | \theta_t, \varphi_t] = \frac{|I|}{n} N_t. \quad (30)$$

Plugging these expectations in (28), we get that,

$$\mathbb{E}[N_{t+1}] = \left(1 - \frac{|I|}{n} (2\eta\epsilon - \eta^2(1 + \epsilon^2)) - \frac{|I|^2}{n^2} (2\eta^2 - \eta^4(1 + \epsilon^2))\right) \mathbb{E}[N_t]. \quad (31)$$

Consequently for $\eta < \epsilon$ we get,

$$\mathbb{E}[N_{t+1}] \geq \left(1 - 2\eta^2 \frac{|I|^2}{n^2} + \eta^2 \frac{|I|}{n}\right) \mathbb{E}[N_t]. \quad (32)$$

To sum-up, if $|I|$ is not large enough (more precisely if $2|I| \leq n$), we have the geometric divergence of the quantity $\mathbb{E}[N_t] := \mathbb{E}[\|\theta_t\|^2 + \|\varphi_t\|^2]$ for any $\eta \geq \epsilon$. \square

C.2 Proof of Theorem 2

Setting of the Proof. We will prove a slightly more general result than Theorem 2. We will work in the context of monotone operator. Let us consider the *general* extrapolation update rule,

$$\begin{cases} \text{Extrapolation:} & \omega_{t+\frac{1}{2}} = \omega_t - \eta_t g_t \\ \text{Update:} & \omega_{t+1} = \omega_t - \eta_t g_{t+1/2}, \end{cases} \quad (33)$$

where g_t depends on ω_t and $g_{t+1/2}$ depends on $\omega_{t+1/2}$. For instance, g_t can either be $F(\omega_t)$, $F_{i_t}(\omega_t)$ or the SVRG estimate defined in (45).

This update rule generalizes (EG) for 2-player games (2P-G) and ExtraSVRG (Alg. 2).

Let us first state a lemma standard in convex analysis (see for instance (Boyd and Vandenberghe, 2004)),

Lemma 1. *Let $\omega \in \Omega$ and $\omega^+ := P_\Omega(\omega + u)$ then for all $\omega' \in \Omega$ we have,*

$$\|\omega^+ - \omega'\|_2^2 \leq \|\omega - \omega'\|_2^2 + 2u^\top (\omega^+ - \omega') - \|\omega^+ - \omega\|_2^2. \quad (34)$$

Proof of Lemma 1. We start by simply developing,

$$\begin{aligned}\|\omega^+ - \omega'\|_2^2 &= \|(\omega^+ - \omega) + (\omega - \omega')\|_2^2 = \|\omega - \omega'\|_2^2 + 2(\omega^+ - \omega)^\top(\omega - \omega') + \|\omega^+ - \omega\|_2^2 \\ &= \|\omega - \omega'\|_2^2 + 2(\omega^+ - \omega)^\top(\omega^+ - \omega') - \|\omega^+ - \omega\|_2^2.\end{aligned}$$

Then since ω^+ is the projection onto the convex set Ω of $\omega + u$ we have that $(\omega^+ - (\omega + u))^\top(\omega^+ - \omega') \leq 0$, $\forall \omega' \in \Omega$, leading to the result of the Lemma. \square

Lemma 2. If F is $(\mu_\theta, \mu_\varphi)$ -strongly monotone for any $\omega, \omega', \omega'' \in \Omega$ we have,

$$\mu_\theta (\|\theta - \theta'\|_2^2 - 2\|\theta' - \theta\|_2^2) + \mu_\varphi (\|\varphi - \varphi'\|_2^2 - 2\|\varphi' - \varphi\|_2^2) \leq 2(F(\omega') - F(\omega''))^\top(\omega' - \omega''), \quad (35)$$

where we noted $\omega := (\theta, \varphi)$.

Proof. By $(\mu_\theta, \mu_\varphi)$ -strong monotonicity,

$$2\mu_\theta \|\theta' - \theta''\|_2^2 + 2\mu_\varphi \|\varphi' - \varphi''\|_2^2 \leq 2(F(\omega'') - F(\omega'))^\top(\omega' - \omega'') \quad (36)$$

and then we use the inequality $2\|\mathbf{a}' - \mathbf{a}''\|_2^2 \geq \|\mathbf{a} - \mathbf{a}''\|_2^2 - 2\|\mathbf{a}' - \mathbf{a}\|_2^2$ to get the result claimed. \square

Using this update rule we can thus deduce the following lemma, the derivation of this lemma is very similar from the derivation of [Harker and Pang \(1990, Lemma 12.1.10\)](#).

Lemma 3. Considering the update rule (33), we have for any $\omega \in \Omega$ and any $t \geq 0$,

$$2\eta_t \mathbf{g}_{t+1/2}^\top(\omega_{t+1/2} - \omega) \leq \|\omega_t - \omega\|_2^2 - \|\omega_{t+1} - \omega\|_2^2 - \|\omega_{t+1/2} - \omega_t\|_2^2 + \eta_t^2 \|\mathbf{g}_t - \mathbf{g}_{t+1/2}\|_2^2. \quad (37)$$

Proof. By applying Lem. 1 for $(\omega, u, \omega^+, \omega') = (\omega_t, -\eta_t \mathbf{g}_{t+1/2}, \omega_{t+1}, \omega)$ and $(\omega, u, \omega^+, \omega') = (\omega_t, -\eta_t \mathbf{g}_t, \omega_{t+1/2}, \omega_{t+1})$, we get,

$$\|\omega_{t+1} - \omega\|_2^2 \leq \|\omega_t - \omega\|_2^2 - 2\eta_t \mathbf{g}_{t+1/2}^\top(\omega_{t+1} - \omega) - \|\omega_{t+1} - \omega_t\|_2^2, \quad (38)$$

and

$$\|\omega_{t+1/2} - \omega_{t+1}\|_2^2 \leq \|\omega_t - \omega_{t+1}\|_2^2 - 2\eta_t \mathbf{g}_t^\top(\omega_{t+1/2} - \omega_{t+1}) - \|\omega_{t+1/2} - \omega_t\|_2^2. \quad (39)$$

Summing (38) and (39) we get,

$$\|\omega_{t+1} - \omega\|_2^2 \leq \|\omega_t - \omega\|_2^2 - 2\eta_t \mathbf{g}_{t+1/2}^\top(\omega_{t+1} - \omega) \quad (40)$$

$$- 2\eta_t \mathbf{g}_t^\top(\omega_{t+1/2} - \omega_{t+1}) - \|\omega_{t+1/2} - \omega_t\|_2^2 - \|\omega_{t+1/2} - \omega_{t+1}\|_2^2 \quad (41)$$

$$\begin{aligned}&= \|\omega_t - \omega\|_2^2 - 2\eta_t \mathbf{g}_{t+1/2}^\top(\omega_{t+1/2} - \omega) - \|\omega_{t+1/2} - \omega_t\|_2^2 - \|\omega_{t+1/2} - \omega_{t+1}\|_2^2 \\ &\quad - 2\eta_t (\mathbf{g}_t - \mathbf{g}_{t+1/2})^\top(\omega_{t+1/2} - \omega_{t+1}).\end{aligned} \quad (42)$$

Then, we can use Young's inequality $-2a^\top b \leq \|a\|_2^2 + \|b\|_2^2$ to get,

$$\begin{aligned}\|\omega_{t+1} - \omega\|_2^2 &\leq \|\omega_t - \omega\|_2^2 - 2\eta_t \mathbf{g}_{t+1/2}^\top(\omega_{t+1/2} - \omega) + \eta_t^2 \|\mathbf{g}_t - \mathbf{g}_{t+1/2}\|_2^2 \\ &\quad + \|\omega_{t+1/2} - \omega_{t+1}\|_2^2 - \|\omega_{t+1/2} - \omega_t\|_2^2 - \|\omega_{t+1/2} - \omega_{t+1}\|_2^2\end{aligned} \quad (43)$$

$$= \|\omega_t - \omega\|_2^2 - 2\eta_t \mathbf{g}_{t+1/2}^\top(\omega_{t+1/2} - \omega) + \eta_t^2 \|\mathbf{g}_t - \mathbf{g}_{t+1/2}\|_2^2 - \|\omega_{t+1/2} - \omega_t\|_2^2. \quad (44)$$

\square

Note that if we would have set $\mathbf{g}_t = \mathbf{0}$ and $\mathbf{g}_{t+1/2}$ any estimate of the gradient at $\boldsymbol{\omega}_t$ we recover the standard lemma for gradient method.

Let us consider *unbiased* estimates of the gradient,

$$\mathbf{g}_i(\boldsymbol{\omega}) := \frac{1}{n\pi_i} (F_i(\boldsymbol{\omega}) - \boldsymbol{\alpha}_i) + \bar{\boldsymbol{\alpha}}, \quad (45)$$

where $\bar{\boldsymbol{\alpha}} := \frac{1}{n} \sum_{j=1}^n \boldsymbol{\alpha}_j$, the index i are (potentially) non-uniformly sampled from $\{1, \dots, n\}$ with replacement according to $\boldsymbol{\pi}$ and $F(\boldsymbol{\omega}) := \frac{1}{n} \sum_{j=1}^n F_j(\boldsymbol{\omega})$. Hence we have that $\mathbb{E}[\mathbf{g}_i(\boldsymbol{\omega})] = F(\boldsymbol{\omega})$, where the expectation is taken with respect to the index i sampled from $\boldsymbol{\pi}$.

We will consider a class of algorithm called *uniform memorization algorithms* first introduced by (Hofmann et al., 2015). This class of algorithms describes a large subset of variance reduced algorithms taking advantage of the finite sum formulation such as SAGA (Defazio et al., 2014), SVRG (Johnson and Zhang, 2013) or q -SAGA and \mathcal{N} -SAGA (Hofmann et al., 2015). In this work, we will use a slightly more general definition of such algorithm in order to be able to handle extrapolation steps:

Definition 9 (Extension of (Hofmann et al., 2015)). *A uniform q -memorization algorithm evolves iterates $(\boldsymbol{\omega}_t)$ according to (33), with \mathbf{g}_t defined in (45) and selecting in each iteration t a random index set J_t of memory locations to update according to,*

$$\boldsymbol{\alpha}_k^{(0)} := F_k(\boldsymbol{\omega}_0), \quad \boldsymbol{\alpha}_k^{(t+1/2)} := \boldsymbol{\alpha}_k^{(t)}, \quad \forall k \in \{1, \dots, n\} \quad \text{and} \quad \boldsymbol{\alpha}_k^{(t+1)} := \begin{cases} F_k(\boldsymbol{\omega}_t) & \text{if } k \in J_t \\ \boldsymbol{\alpha}_k^{(t)} & \text{otherwise.} \end{cases} \quad (46)$$

such that any k has the same probability q/n to be updated, i.e., $P\{k\} = \sum_{J_t, k \in J_t} P(J_t) = q/n$, $\forall k \in \{1, \dots, n\}$.

In the case of SVRG, either $J_t = \emptyset$ or $J_t = \{1, \dots, n\}$ (when we update the snapshot).

We have the following lemmas,

Lemma 4. *For any $t \geq 0$, if we consider a q -memorization algorithm we have*

$$\mathbb{E}[\|\mathbf{g}_t - \mathbf{g}_{t+1/2}\|^2] \leq 10\mathbb{E}[\|\frac{1}{n\pi_i}(F_i(\boldsymbol{\omega}^*) - \boldsymbol{\alpha}_i^{(t)})\|^2] + 10\mathbb{E}[\|\frac{1}{n\pi_i}(F_i(\boldsymbol{\omega}^*) - F_i(\boldsymbol{\omega}_t))\|^2] + 5\bar{L}^2\mathbb{E}[\|\boldsymbol{\omega}_t - \boldsymbol{\omega}_{t+1/2}\|^2].$$

Proof. We use an extended version of Young's inequality: $\|\sum_{i=1}^k \mathbf{a}_i\|^2 \leq k \sum_{i=1}^k \|\mathbf{a}_i\|^2$,

$$\begin{aligned} \left\| \sum_{i=1}^k \mathbf{a}_i \right\|^2 &= \sum_{i,j=1}^k \mathbf{a}_i^\top \mathbf{a}_j \\ &\leq \frac{1}{2} \sum_{i,j=1}^k \|\mathbf{a}_i\|^2 + \|\mathbf{a}_j\|^2 \\ &= k \sum_{i=1}^k \|\mathbf{a}_i\|^2, \end{aligned}$$

where we used that $2\mathbf{a}^\top \mathbf{b} \leq \|\mathbf{a}\|^2 + \|\mathbf{b}\|^2$. We combine Young's inequality with the definition of q -memorization algorithm: $\mathbf{g}_t = \frac{1}{n\pi_i}(F_i(\boldsymbol{\omega}_t) - \bar{\boldsymbol{\alpha}}_i^{(t)})$ and $\mathbf{g}_{t+1/2} = \frac{1}{n\pi_j}(F_j(\boldsymbol{\omega}_{t+1/2}) - \bar{\boldsymbol{\alpha}}_j^{(t)})$ to get (we omit

the t subscript for i and j and we note $\bar{\alpha}_i^{(t)} := \alpha_i^{(t)} - n\pi_i \bar{\alpha}^{(t)}$,

$$\begin{aligned}
\|g_t - g_{t+1/2}\|^2 &= \left\| \frac{1}{n\pi_i} (F_i(\omega_t) - \bar{\alpha}_i^{(t)}) - \frac{1}{n\pi_j} (F_j(\omega_{t+1/2}) - \bar{\alpha}_j^{(t)}) \right\|^2 \\
&= \left\| \frac{1}{n\pi_i} (F_i(\omega_t) - \bar{\alpha}_i^{(t)}) + \frac{1}{n\pi_j} (F_j(\omega_t) - F_j(\omega_{t+1/2})) + \frac{1}{n\pi_j} (\bar{\alpha}_j^{(t)} - F_j(\omega_t)) \right\|^2 \\
&\leq 5\mathbb{E}[\left\| \frac{1}{n\pi_i} (F_i(\omega^*) - \bar{\alpha}_i^{(t)}) \right\|^2] + 5\mathbb{E}[\left\| \frac{1}{n\pi_j} (\bar{\alpha}_j^{(t)} - F_j(\omega^*)) \right\|^2] \\
&\quad + 5\mathbb{E}[\left\| \frac{1}{n\pi_i} (F_i(\omega^*) - F_i(\omega_t)) \right\|^2] + 5\mathbb{E}[\left\| \frac{1}{n\pi_j} (F_j(\omega^*) - F_j(\omega_t)) \right\|^2] \\
&\quad + 5\mathbb{E}[\left\| \frac{1}{n\pi_j} (F_j(\omega_t) - F_j(\omega_{t+1/2})) \right\|^2]
\end{aligned}$$

Notice that since i_t and j_t are independently sampled from the same distribution we have

$$\mathbb{E}[\left\| \frac{1}{n^2\pi_{j_t}^2} \|F_{j_t}(\omega^*) - \alpha_{j_t}^{(t)}\|^2\right\|] = \mathbb{E}[\left\| \frac{1}{n^2\pi_{i_t}^2} \|F_{i_t}(\omega^*) - \alpha_{i_t}^{(t)}\|^2\right\|]. \quad (47)$$

Note that we have (using that $\mathbb{E}[F_i(\omega^*)] = 0$ and $\mathbb{E}[\alpha_i^{(t)}] = \bar{\alpha}^{(t)}$),

$$\mathbb{E}[\left\| \frac{1}{n\pi_i} (F_i(\omega^*) - \bar{\alpha}_i^{(t)}) \right\|^2] = \mathbb{E}[\left\| \frac{1}{n\pi_i} (F_i(\omega^*) - \alpha_{i_t}^{(t)}) \right\|^2] - \|\bar{\alpha}^{(t)}\|^2 \leq \mathbb{E}[\left\| \frac{1}{n\pi_i} (F_i(\omega^*) - \alpha_{i_t}^{(t)}) \right\|^2] \quad (48)$$

By assuming that each F_i is L_i -Lipschitz we get,

$$\mathbb{E}[\left\| \frac{1}{n^2\pi_{j_t}^2} \|F_j(\omega_t) - F_j(\omega_{t+1/2})\|^2\right\|] = \frac{1}{n^2} \sum_{j=1}^n \frac{1}{\pi_j} \mathbb{E}[\left\| F_j(\omega_t) - F_j(\omega_{t+1/2}) \right\|^2] \quad (49)$$

$$\leq \frac{1}{n^2} \sum_{j=1}^n \frac{L_j^2}{\pi_j} \mathbb{E}[\|\omega_t - \omega_{t+1/2}\|^2] \quad (50)$$

$$= \bar{L}^2 \mathbb{E}[\|\omega_t - \omega_{t+1/2}\|^2], \quad (51)$$

where $\bar{L}^2 := \frac{1}{n^2} \sum_{i=1}^n \frac{L_i^2}{\pi_i}$. Note that ω_t and $\omega_{t+1/2}$ do not depend on j_t (which is the index sampled for the update step), that is not the case for i (the index for the extrapolation step) since $\omega_{t+1/2}$ is the result of the extrapolation. \square

This lemma make appear the quantity $\mathbb{E}[\left\| \frac{1}{n\pi_i} (F_i(\omega^*) - \bar{\alpha}_i^{(t)}) \right\|^2]$ that we need to bound. In order to do that we prove the following lemma,

Lemma 5. *Let $(\alpha_j^{(t)})$ be updated according to the rules of a q -uniform memorization algorithm (Def. 9). Let us note $H_t := \frac{1}{n} \sum_{i=1}^n \frac{1}{n\pi_i} \|F_i(\omega^*) - \alpha_i^{(t)}\|^2$. For any $t \in \mathbb{N}$,*

$$\mathbb{E}[H_{t+1}] = \frac{q}{n} \mathbb{E}[\left\| \frac{1}{n\pi_{i_t}} (F_{i_t}(\omega_t) - F_{i_t}(\omega^*)) \right\|^2] + \frac{n-q}{n} \mathbb{E}[H_t]. \quad (52)$$

Proof. We will use the definition of q -uniform memorization algorithms (saying that α_i is updated at time $t + 1$ with probability q/n). We call this event " i updated",

$$\begin{aligned}
\mathbb{E}[H_{t+1}] &:= \mathbb{E}\left[\frac{1}{n} \sum_{i=1}^n \frac{1}{n\pi_i} \|\alpha_i^{(t+1)} - F_i(\omega^*)\|^2\right] \\
&= \frac{1}{n} \mathbb{E}\left[\sum_{i \text{ updated}} \frac{1}{n\pi_i} \|\alpha_i^{(t+1)} - F_i(\omega^*)\|^2 + \sum_{i \text{ not updated}} \frac{1}{n\pi_i} \|\alpha_i^{(t+1)} - F_i(\omega^*)\|^2\right] \\
&= \frac{1}{n} \mathbb{E}\left[\sum_{i \text{ updated}} \frac{1}{n\pi_i} \|F_i(\omega_t) - F_i(\omega^*)\|^2 + \sum_{i \text{ not updated}} \frac{1}{n\pi_i} \|\alpha_i^{(t)} - F_i(\omega^*)\|^2\right] \\
&= \frac{1}{n} \sum_{i=1}^n \mathbf{P}(i \text{ updated}) \frac{1}{n\pi_i} \mathbb{E}[\|F_i(\omega_t) - F_i(\omega^*)\|^2] + \frac{1}{n} \sum_{i=1}^n \mathbf{P}(i \text{ not updated}) \frac{1}{n\pi_i} \mathbb{E}[\|\alpha_i^{(t)} - F_i(\omega^*)\|^2] \\
&= \frac{q}{n} \mathbb{E}[\|\frac{1}{n\pi_{i_t}} (F_{i_t}(\omega_t) - F_{i_t}(\omega^*))\|^2] + \frac{n-q}{n} \mathbb{E}[H_t]
\end{aligned}$$

□

Using all these lemmas we can prove our theorem.

Theorem' 2. Under Assumption 1, after t iterations, the iterate ω_t computed by a q -memorization algorithm with step-sizes $(\eta_\theta, \eta_\phi) \leq ((40\bar{\ell}_\theta)^{-1}, (40\bar{\ell}_\phi)^{-1})$ verifies:

$$\mathbb{E}[\|\omega_t - \omega^*\|_2^2] \leq \left(1 - \min\left\{\frac{\eta\mu}{2} + \frac{9\eta^2\gamma^2}{10}, \frac{2q}{5n}\right\}\right)^t \mathbb{E}[\|\omega_0 - \omega^*\|_2^2]. \quad (53)$$

Proof. In this proof we will consider a constant step-size $\eta_t = (\eta_\theta, \eta_\phi)$. For simplicity of notations we will consider the notation,

$$\begin{aligned}
\bar{L}^2 \|\omega\|^2 &:= \bar{L}_\theta^2 \|\theta\|^2 + \bar{L}_\phi^2 \|\phi\|^2, \quad \eta^2 \|\omega\|^2 := \eta_\theta^2 \|\theta\|^2 + \eta_\phi^2 \|\phi\|^2, \quad \mu \|\omega\|^2 := \mu_\theta^2 \|\theta\|^2 + \mu_\phi^2 \|\phi\|^2 \\
\eta\mu &= (\eta_\theta \mu_\theta, \eta_\phi \mu_\phi), \quad \sigma \bar{L}^2 = (\sigma_\theta \bar{L}_\theta^2, \sigma_\phi \bar{L}_\phi^2) \quad \text{and} \quad \eta^2 \bar{L}^2 = (\eta_\theta^2 \bar{L}_\theta^2, \eta_\phi^2 \bar{L}_\phi^2).
\end{aligned}$$

We start by recalling Lemma 3,

$$\|\omega_{t+1} - \omega^*\|_2^2 \leq \|\omega_t - \omega^*\|_2^2 - 2\eta \mathbf{g}_{t+1/2}^\top (\omega_{t+1/2} - \omega^*) - (1 - 2\eta\mu) \|\omega_{t+1/2} - \omega_t\|_2^2 + \eta^2 \|\mathbf{g}_t - \mathbf{g}_{t+1/2}\|_2^2. \quad (54)$$

We can then take the expectation and plug-in the expression of $\mathbb{E}[\|\mathbf{g}_t - \mathbf{g}_{t+1/2}\|_2^2]$ from Lemma 4,

$$\begin{aligned}
\mathbb{E}[\|\omega_{t+1} - \omega^*\|_2^2] &\leq \mathbb{E}[\|\omega_t - \omega^*\|_2^2] - 2\eta \mathbb{E}[F(\omega_{t+1/2})^\top (\omega_{t+1/2} - \omega^*)] - (1 - 2\eta\mu - 5\eta^2 \bar{L}^2) \mathbb{E}[\|\omega_{t+1/2} - \omega_t\|_2^2] \\
&\quad + \eta^2 (10\mathbb{E}[\|\frac{1}{n\pi_i} (F_i(\omega^*) - \alpha_i^{(t)})\|^2] + 10\mathbb{E}[\|\frac{1}{n\pi_i} (F_i(\omega^*) - F_i(\omega_t))\|^2]).
\end{aligned}$$

Let us define $\mathcal{L}_t := \mathbb{E}[\|\omega_t - \omega^*\|_2^2] + \sigma \mathbb{E}[H_t]$, where $H_t := \frac{1}{n} \sum_{i=1}^n \frac{1}{n\pi_i} \|F_i(\omega^*) - \alpha_i^{(t)}\|^2$. We can combine (54) with Lemma 5 multiplied by a constant $\sigma > 0$ that we will set later to get

$$\begin{aligned}
\mathcal{L}_{t+1} &= \mathbb{E}[\|\omega_{t+1} - \omega^*\|_2^2] + \sigma \mathbb{E}[H_{t+1}] \\
&\leq \mathbb{E}[\|\omega_t - \omega^*\|_2^2] - 2\eta \mathbb{E}[F(\omega_{t+1/2})^\top (\omega_{t+1/2} - \omega^*)] - (1 - 2\eta\mu - 5\eta^2 \bar{L}^2) \mathbb{E}[\|\omega_{t+1/2} - \omega_t\|_2^2] \\
&\quad + (\frac{\sigma q}{n} + 10\eta^2) \mathbb{E}[\|\frac{1}{n\pi_i} (F_i(\omega^*) - F_i(\omega_t))\|^2] + (\frac{10\eta^2}{\sigma} + \frac{n-q}{n}) \sigma \mathbb{E}[H_t].
\end{aligned}$$

Since i_t and j_t are independently drawn from the same distribution, we have, $\mathbb{E}[\|\frac{1}{n\pi_i}(F_i(\boldsymbol{\omega}^*) - F_i(\boldsymbol{\omega}_t))\|^2] = \mathbb{E}[\|\frac{1}{n\pi_j}(F_j(\boldsymbol{\omega}^*) - F_j(\boldsymbol{\omega}_t))\|^2]$ and thus,

$$\begin{aligned}
\mathcal{L}_{t+1} &\leq \mathbb{E}[\|\boldsymbol{\omega}_t - \boldsymbol{\omega}^*\|_2^2] - 2\eta\mathbb{E}[F(\boldsymbol{\omega}_{t+1/2})^\top(\boldsymbol{\omega}_{t+1/2} - \boldsymbol{\omega}^*)] - (1 - 2\eta\mu - 5\eta^2\bar{L}^2)\|\boldsymbol{\omega}_{t+1/2} - \boldsymbol{\omega}_t\|_2^2 \\
&\quad + (\frac{\sigma q}{n} + 10\eta^2)\mathbb{E}[\|\frac{1}{n\pi_j}(F_j(\boldsymbol{\omega}^*) - F_j(\boldsymbol{\omega}_t))\|^2] + (\frac{10\eta^2}{\sigma} + \frac{n-q}{n})\sigma\mathbb{E}[H_t] \\
&\leq \mathbb{E}[\|\boldsymbol{\omega}_t - \boldsymbol{\omega}^*\|_2^2] - (1 - 2\eta\mu - 5\eta^2\bar{L}^2 - 2(\frac{\sigma q}{n} + 10\eta^2)\bar{L}^2)\|\boldsymbol{\omega}_{t+1/2} - \boldsymbol{\omega}_t\|_2^2 \\
&\quad - 2\eta\mathbb{E}[F(\boldsymbol{\omega}_{t+1/2})^\top(\boldsymbol{\omega}_{t+1/2} - \boldsymbol{\omega}^*)] + 2(\frac{\sigma q}{n} + 10\eta^2)\mathbb{E}[\|\frac{1}{n\pi_j}(F_j(\boldsymbol{\omega}^*) - F_j(\boldsymbol{\omega}_{t+1/2}))\|^2] \\
&\quad + (\frac{10\eta^2}{\sigma} + \frac{n-q}{n})\sigma\mathbb{E}[H_t] \\
&\leq \mathbb{E}[\|\boldsymbol{\omega}_t - \boldsymbol{\omega}^*\|_2^2] - (1 - 2\eta\mu - 5\eta^2\bar{L}^2 - 2(\frac{\sigma q}{n} + 10\eta^2)\bar{L}^2)\|\boldsymbol{\omega}_{t+1/2} - \boldsymbol{\omega}_t\|_2^2 \\
&\quad - 2\eta\mathbb{E}[F(\boldsymbol{\omega}_{t+1/2})^\top(\boldsymbol{\omega}_{t+1/2} - \boldsymbol{\omega}^*)] + (\frac{10\eta^2}{\sigma} + \frac{n-q}{n})\sigma\mathbb{E}[H_t] \\
&\quad + 2(\frac{\sigma q}{n} + 10\eta^2)\mathbb{E}[\frac{\ell_j}{n^2\pi_j^2}(F_j(\boldsymbol{\omega}^*) - F_j(\boldsymbol{\omega}_{t+1/2}))^\top(\boldsymbol{\omega}^* - \boldsymbol{\omega}_{t+1/2})]
\end{aligned}$$

where for the second inequality we used Young's inequality and the Lipchitzness of F_j and for the last one we used the co-coercivity of F_j :

$$\|F_j(\boldsymbol{\omega}) - F_j(\boldsymbol{\omega}')\|^2 \leq \ell_j(F_j(\boldsymbol{\omega}') - F_j(\boldsymbol{\omega}))^\top(\boldsymbol{\omega}' - \boldsymbol{\omega}). \quad (55)$$

Thus using $\pi_j = \frac{\ell_j}{\sum_j \ell_j}$, we get

$$\begin{aligned}
\mathcal{L}_{t+1} &\leq \mathbb{E}[\|\boldsymbol{\omega}_t - \boldsymbol{\omega}^*\|_2^2] - (1 - 2\eta\mu - 5\eta^2\bar{L}^2 - 2(\frac{\sigma q}{n} + 10\eta^2)\bar{L}^2)\|\boldsymbol{\omega}_{t+1/2} - \boldsymbol{\omega}_t\|_2^2 \\
&\quad - 2\eta\mathbb{E}[F(\boldsymbol{\omega}_{t+1/2})^\top(\boldsymbol{\omega}_{t+1/2} - \boldsymbol{\omega}^*)] + 2(\frac{10\eta^2}{\sigma} + \frac{n-q}{n})\sigma\mathbb{E}[H_t] \\
&\quad + 2\bar{\ell}(\frac{\sigma q}{n} + 10\eta^2)\mathbb{E}[\frac{1}{n\pi_j}(F_j(\boldsymbol{\omega}^*) - F_j(\boldsymbol{\omega}_{t+1/2}))^\top(\boldsymbol{\omega}^* - \boldsymbol{\omega}_{t+1/2})] \\
&= \mathbb{E}[\|\boldsymbol{\omega}_t - \boldsymbol{\omega}^*\|_2^2] - (1 - 2\eta\mu - 5\eta^2\bar{L}^2 - 2(\frac{\sigma q}{n} + 10\eta^2)\bar{L}^2)\|\boldsymbol{\omega}_{t+1/2} - \boldsymbol{\omega}_t\|_2^2 \\
&\quad - 2\eta\mathbb{E}[F(\boldsymbol{\omega}_{t+1/2})^\top(\boldsymbol{\omega}_{t+1/2} - \boldsymbol{\omega}^*)] + 2\bar{\ell}(\frac{\sigma q}{n} + 10\eta^2)\mathbb{E}[F(\boldsymbol{\omega}_{t+1/2})^\top(\boldsymbol{\omega}_{t+1/2} - \boldsymbol{\omega}^*)] \\
&\quad + (\frac{10\eta^2}{\sigma} + \frac{n-q}{n})\sigma\mathbb{E}[H_t]
\end{aligned}$$

where $\bar{\ell} := \frac{1}{n} \sum_i \ell_i$. Now we can set $\frac{20\eta^2}{\sigma} = \frac{q}{n}$ to get

$$\begin{aligned}
\mathcal{L}_{t+1} &\leq \mathbb{E}[\|\boldsymbol{\omega}_t - \boldsymbol{\omega}^*\|_2^2] - (1 - 2\eta\mu - 65\eta^2\bar{L}^2)\|\boldsymbol{\omega}_{t+1/2} - \boldsymbol{\omega}_t\|_2^2 \\
&\quad - \eta(2 - 60\bar{\ell}\eta)\mathbb{E}[F(\boldsymbol{\omega}_{t+1/2})^\top(\boldsymbol{\omega}_{t+1/2} - \boldsymbol{\omega}^*)] + (1 - \frac{q}{2n})\sigma\mathbb{E}[H_t].
\end{aligned}$$

Finally with $\eta \leq \frac{1}{40\bar{\ell}}$ (note that we always have $\bar{\ell} \geq \bar{L}$ because $\ell_i \geq L_i$) we get

$$\mathcal{L}_{t+1} \leq \mathbb{E}[\|\boldsymbol{\omega}_t - \boldsymbol{\omega}^*\|_2^2] - \eta\frac{1}{2}\mathbb{E}[F(\boldsymbol{\omega}_{t+1/2})^\top(\boldsymbol{\omega}_{t+1/2} - \boldsymbol{\omega}^*)] - \frac{9}{10}\mathbb{E}[\|\boldsymbol{\omega}_{t+1/2} - \boldsymbol{\omega}_t\|_2^2] + (1 - \frac{q}{2n})\sigma\mathbb{E}[H_t].$$

We finally use the projection-type error bound $\|F_i(\boldsymbol{\omega}_t) - F_i(\boldsymbol{\omega}^*)\|^2 \geq \gamma_i^2\|\boldsymbol{\omega}_t - \boldsymbol{\omega}^*\|^2$ the same way as (Azizian et al., 2019) to get,

$$\begin{aligned}
\|\boldsymbol{\omega}_{t+1/2} - \boldsymbol{\omega}_t\|_2^2 &= \eta^2\|\frac{1}{n\pi_i}(F_i(\boldsymbol{\omega}_t) - \bar{\boldsymbol{\alpha}}_i^{(t)})\|^2 \\
&\geq \frac{\eta^2}{2}\|\frac{1}{n\pi_i}(F_i(\boldsymbol{\omega}_t) - F_i(\boldsymbol{\omega}^*))\|^2 - \eta^2\|\frac{1}{n\pi_i}(F_i(\boldsymbol{\omega}^*) - \bar{\boldsymbol{\alpha}}_i^{(t)})\|^2 \\
&\geq \frac{\gamma_i^2\eta^2}{2}\|\frac{1}{n\pi_i}(\boldsymbol{\omega}_t - \boldsymbol{\omega}^*)\|^2 - \eta^2\|\frac{1}{n\pi_i}(F_i(\boldsymbol{\omega}^*) - \bar{\boldsymbol{\alpha}}_i^{(t)})\|^2.
\end{aligned}$$

Thus we have that,

$$\mathcal{L}_{t+1} \leq \mathbb{E}[\|\omega_t - \omega^*\|_2^2] - \eta \frac{1}{2} \mathbb{E}[F(\omega_{t+1/2})^\top (\omega_{t+1/2} - \omega^*)] - \frac{\bar{\gamma}^2 \eta^2}{2} \mathbb{E}[\|\omega_t - \omega^*\|_2^2] + (1 - \frac{q}{2n} + \frac{9q}{100n}) \sigma \mathbb{E}[H_t],$$

where $\bar{\gamma}^2 := \frac{1}{n} \sum_{k=1}^n \frac{\gamma_k^2}{n\pi_k}$. We can thus conclude the proof using the strong convexity of F ,

$$\mathcal{L}_{t+1} \leq \left(1 - \min \left\{ \left(\frac{\eta\mu}{2} + \frac{9\eta^2\bar{\gamma}^2}{20} \right), \frac{2q}{5n} \right\} \right) \mathcal{L}_t.$$

□

D Details on the SVRE-GAN Algorithm

D.1 Practical Aspect

Noise dataset. Variance reduction is usually performed on finite sum dataset. However, the noise dataset in GANs (sampling from the noise variable z for the generator G) is in practice considered as an infinite dataset. We considered several ways to cope with this:

- Infinitely taking new samples from a predefined latent distribution p_g . In this case, from a theoretical point of view, in terms of using finite sum formulation, there is no convergence guarantee for SVRE even in the strongly convex case. Moreover, the estimators (63) and (64) are biased estimator of the gradient (as μ_D and μ_G do not estimate the full expectation but a finite sum).
- Sampling a different noise dataset at each epoch, i.e. considering a different finite sum at each epoch. In that case, we are performing a variance reduction of this finite sum over the epoch.
- Fix a finite sum noise dataset for the entire training.

In practice, we did not notice any notable difference between the three alternatives.

Adaptive methods. Particular choices such as the optimization method (*e.g.* Adam (Kingma and Ba, 2015)), learning rates, and normalization, have been established in practice as almost *prerequisite* for convergence⁷, in contrast to supervised classification problems where they have been shown to only provide a marginal value (Wilson et al., 2017). To our knowledge, SVRE is the only method that works with a constant step size for GANs on non-trivial datasets. This combined with the fact that recent works empirically tune the first moment controlling hyperparameter to 0 (β_1 , see below) and the variance reduction (VR) one (β_2 , see below) to a non-zero value, sheds light on the reason behind the success of Adam on GANs.

However, combining SVRE with adaptive step size scheme on GANs remains an open problem. We first briefly describe the update rule of Adam, and then we propose a new adaptation of it that is more suitable for VR methods, which we refer to as variance reduced Adam (VRAd).

Adam. Adam stores an exponentially decaying average of both past gradients m_t and squared gradients v_t , for each parameter of the model:

$$m_t = \beta_1 m_{t-1} + (1 - \beta_1) g_t \tag{56}$$

$$v_t = \beta_2 v_{t-1} + (1 - \beta_2) g_t^2, \tag{57}$$

⁷For instance, Daskalakis et al. (2018); Gidel et al. (2019a) plugged Adam into their principled method to get better results.

Algorithm 2 Pseudocode for SVRE-GAN.

```

1: Input: dataset  $\mathcal{D}$ , noise dataset  $\mathcal{Z}$  ( $|\mathcal{Z}| = |\mathcal{D}| = n$ ), stopping iteration  $T$ , learning rates  $\eta_D, \eta_G$ , generator
   loss  $\mathcal{L}^G$ , discriminator loss  $\mathcal{L}^D$ , mini-batch size  $B$ .
2: Initialize:  $D, G$ 
3: for  $e = 0$  to  $T-1$  do
4:    $D^S = D$  and  $\mu_D = \frac{1}{n} \sum_{i=1}^n \sum_{j=1}^n \nabla_D \mathcal{L}^D(G^S, D^S, \mathcal{D}_j, \mathcal{Z}_i)$ 
5:    $G^S = G$  and  $\mu_G = \frac{1}{n} \sum_{i=1}^n \nabla_G \mathcal{L}^G(G^S, D^S, \mathcal{Z}_i)$ 
6:    $N \sim \text{Geom}(B/n)$  (length of the epoch)
7:   for  $i = 0$  to  $N-1$  do
8:     Sample mini-batches  $(n_d, n_z)$ ; do extrapolation:
9:      $\tilde{D} = D - \eta_D \mathbf{d}_D(G, D, G^S, D^S, n_z)$   $\triangleright$  (63)
10:     $\tilde{G} = G - \eta_G \mathbf{d}_G(G, D, G^S, D^S, n_d, n_z)$   $\triangleright$  (64)
11:    Sample new mini-batches  $(n_d, n_z)$ ; do update:
12:     $D = D - \eta_D \mathbf{d}_D(\tilde{G}, \tilde{D}, G^S, D^S, n_z)$   $\triangleright$  (63)
13:     $G = G - \eta_G \mathbf{d}_G(\tilde{G}, \tilde{D}, G^S, D^S, n_d, n_z)$   $\triangleright$  (64)
14: Output:  $G, D$ 

```

where $\beta_1, \beta_2 \in [0, 1]$, $m_0 = 0$, $v_0 = 0$, and $t = 1, \dots, T$ denotes the iteration. m_t and v_t are respectively the estimates of the first and the second moments of the stochastic gradient. To compensate the bias toward 0 due to initialization, Kingma and Ba (2015) propose to use bias-corrected estimates of these first two moments:

$$\hat{m}_t = \frac{m_t}{1 - \beta_1^t} \quad (58)$$

$$\hat{v}_t = \frac{v_t}{1 - \beta_2^t}. \quad (59)$$

The Adam update rule can be described as:

$$\omega_{t+1} = \omega_t - \eta \frac{\hat{m}_t}{\sqrt{\hat{v}_t} + \epsilon}. \quad (60)$$

Adam can be understood as an approximate gradient method with a diagonal step size of $\eta_{Adam} := \frac{\eta}{\sqrt{v_t} + \epsilon}$. Since VR methods aim to provide a vanishing v_t , they lead to a too large step-size η_{Adam} of $\frac{\eta}{\epsilon}$. This could indicate that the update rule of Adam may not be a well-suited method to combine with VR methods.

VRAd. This motivates the introduction of a new Adam-inspired variant of adaptive step sizes that maintain a reasonable size even when v_t vanishes,

$$\omega_{t+1} = \omega_t - \eta \frac{|\hat{m}_t|}{\sqrt{\hat{v}_t} + \epsilon} \hat{m}_t. \quad (\text{VRAd})$$

This adaptive variant of Adam is motivated by the step size $\eta^* = \eta \frac{m_t^2}{v_t}$ derived by Schaul et al. (2013). (VRAd) is simply the square-root of η^* in order to stick with Adam's scaling of v_t .

D.2 SVRE-GAN

In order to cope with the issues introduced by the stochastic game formulation of the GAN models, we proposed the SVRE algorithm Alg. 1 which combines SVRG and extragradient method. We refer to

the method of applying SVRE to train GANs as the *SVRE-GAN* method, and we describe it in detail in Alg. 2 (generalizing it with mini-batching, but using uniform probabilities). Assuming that we have $\mathcal{D}[n_d]$ and $\mathcal{Z}[n_z]$, respectively two mini-batches of size B of the true dataset and the noise dataset, we compute $\nabla_D \mathcal{L}^D(G, D, \mathcal{D}[n_d], \mathcal{Z}[n_z])$ and $\nabla_G \mathcal{L}^G(G, D, \mathcal{Z}[n_z])$ the respective mini-batches gradient of the discriminator and the generator:

$$\nabla_D \mathcal{L}^D(G, D, \mathcal{D}[n_d], \mathcal{Z}[n_z]) := \frac{1}{|n_z|} \frac{1}{|n_d|} \sum_{i \in n_z} \sum_{j \in n_d} \nabla_D \mathcal{L}^D(G, D, \mathcal{D}_j, \mathcal{Z}_i) \quad (61)$$

$$\nabla_G \mathcal{L}^G(G, D, \mathcal{Z}[n_z]) := \frac{1}{|n_z|} \sum_{i \in n_z} \nabla_G \mathcal{L}^G(G, D, \mathcal{Z}_i), \quad (62)$$

where \mathcal{Z}_i and \mathcal{D}_j are respectively the i^{th} example of the noise dataset and the j^{th} of the true dataset. Note that n_z and n_d are lists and thus that we allow repetitions in the summations over n_z and n_d . The variance reduced gradient of the SVRG method are thus given by:

$$\mathbf{d}_D(G, D, G^S, D^S) := \boldsymbol{\mu}_D + \nabla_D \mathcal{L}^D(G, D, \mathcal{D}[n_d], \mathcal{Z}[n_z]) - \nabla_D \mathcal{L}^D(G^S, D^S, \mathcal{D}[n_d], \mathcal{Z}[n_z]) \quad (63)$$

$$\mathbf{d}_G(G, D, G^S, D^S) := \boldsymbol{\mu}_G + \nabla_G \mathcal{L}^G(G, D, \mathcal{Z}[n_z]) - \nabla \mathcal{L}^G(G^S, D^S, \mathcal{Z}[n_z]), \quad (64)$$

where G^S and D^S are the snapshots and $\boldsymbol{\mu}_D$ and $\boldsymbol{\mu}_G$ their respective gradients.

Alg. 2 summarizes the SVRG optimization extended to GAN. To obtain that $\mathbb{E} [\nabla_{\boldsymbol{\theta}^S} \mathcal{L}(\boldsymbol{\theta}^S, \boldsymbol{\varphi}^S, \cdot) - \boldsymbol{\mu}]$ vanishes, when updating $\boldsymbol{\theta}$ and $\boldsymbol{\varphi}$ where the expectation is over samples of \mathcal{D} and \mathcal{Z} respectively, we use the snapshot networks $\boldsymbol{\theta}^S$ and $\boldsymbol{\varphi}^S$ for the second term in lines 9, 10, 12 and 13. Moreover, the noise dataset $\mathcal{Z} \sim p_z$, where $|\mathcal{Z}| = |\mathcal{D}| = n$, is fixed. Empirically we observe that directly sampling from p_z (contrary to fixing the noise dataset and re-sampling it with frequency m) does not impact the performance, as $|\mathcal{Z}|$ is usually high.

Note that the double sum in Line 4 can be written as two sums because of the separability of the expectations in typical GAN objectives. Thus the time complexity for calculating $\boldsymbol{\mu}^D$ is still $O(n)$ and not $O(n^2)$ which would be prohibitively expensive.

E Restarted SVRE

Alg. 3 describes the restarted version of SVRE presented in § 3.3. With a probability p (fixed) before the computation of $\boldsymbol{\mu}_{\boldsymbol{\varphi}}^S$ and $\boldsymbol{\mu}_{\boldsymbol{\theta}}^S$, we decide whether to restart SVRE (by using the averaged iterate as the new starting point—Alg. 3, Line 6— $\bar{\boldsymbol{\omega}}_t$) or computing the batch snapshot at a point $\boldsymbol{\omega}_t$.

F Details on the implementation

For our experiments, we used the PyTorch⁸ deep learning framework, whereas for computing the FID and IS metrics, we used the provided implementations in Tensorflow⁹.

F.1 Metrics

We provide more details about the metrics enumerated in § 4. Both FID and IS use: (i) the *Inception v3 network* (Szegedy et al., 2015) that has been trained on the ImageNet dataset consisting of ~ 1 million RGB images of 1000 classes, $C = 1000$. (ii) a sample of m generated images $x \sim p_g$, where usually $m = 50000$.

⁸<https://pytorch.org/>

⁹<https://www.tensorflow.org/>

Algorithm 3 Pseudocode for Restarted SVRE.

```
1: Input: Stopping time  $T$ , learning rates  $\eta_\theta, \eta_\varphi$ , both players' losses  $\mathcal{L}^G$  and  $\mathcal{L}^D$ , probability of restart  $p$ .
2: Initialize:  $\varphi, \theta, t = 0$  ▷  $t$  is for the online average computation.
3: for  $e = 0$  to  $T-1$  do
4:   Draw  $\text{restart} \sim B(p)$ . ▷ Check if we restart the algorithm.
5:   if  $\text{restart}$  and  $e > 0$  then
6:      $\varphi \leftarrow \bar{\varphi}, \theta \leftarrow \bar{\theta}$  and  $t = 1$ 
7:      $\varphi^S \leftarrow \varphi$  and  $\mu_\varphi^S \leftarrow \frac{1}{|\mathcal{Z}|} \sum_{i=1}^n \nabla_\varphi \mathcal{L}_i^D(\theta^S, \varphi^S)$ 
8:      $\theta^S \leftarrow \theta$  and  $\mu_\theta^S \leftarrow \frac{1}{|\varphi|} \sum_{i=1}^n \nabla_\theta \mathcal{L}_i^G(\theta^S, \varphi^S)$ 
9:      $N \sim \text{Geom}(1/n)$  ▷ Length of the epoch.
10:    for  $i = 0$  to  $N-1$  do
11:      Sample  $i_\theta \sim \pi_\theta, i_\varphi \sim \pi_\varphi$ , do extrapolation:
12:       $\tilde{\varphi} \leftarrow \varphi - \eta_\theta d_\varphi(\theta, \varphi, \theta^S, \varphi^S), \tilde{\theta} \leftarrow \theta - \eta_\varphi d_\theta(\theta, \varphi, \theta^S, \varphi^S)$  ▷ (63) and (64)
13:      Sample  $i_\theta \sim \pi_\theta, i_\varphi \sim \pi_\varphi$ , do update:
14:       $\varphi \leftarrow \varphi - \eta_\theta d_\varphi(\tilde{\theta}, \tilde{\varphi}, \theta^S, \varphi^S), \theta \leftarrow \theta - \eta_\varphi d_\theta(\tilde{\theta}, \tilde{\varphi}, \theta^S, \varphi^S)$  ▷ (63) and (64)
15:       $\bar{\theta} \leftarrow \frac{t}{t+1} \bar{\theta} + \frac{1}{t+1} \theta$  and  $\bar{\varphi} \leftarrow \frac{t}{t+1} \bar{\varphi} + \frac{1}{t+1} \varphi$  ▷ Online computation of the average.
16:       $t \leftarrow t + 1$  ▷ Increment  $t$  for the online average computation.
17: Output:  $\theta, \varphi$ 
```

F.1.1 Inception Score

Given an image x , IS uses the softmax output of the Inception network $p(y|x)$ which represents the probability that x is of class $c_i, i \in 1 \dots C$, i.e., $p(y|x) \in [0, 1]^C$. It then computes the marginal class distribution $p(y) = \int_x p(y|x) p_g(x)$. IS measures the Kullback–Leibler divergence \mathbb{D}_{KL} between the predicted conditional label distribution $p(y|x)$ and the marginal class distribution $p(y)$. More precisely, it is computed as follows:

$$IS(G) = \exp(\mathbb{E}_{x \sim p_g}[\mathbb{D}_{KL}(p(y|x) || p(y))]) = \exp\left(\frac{1}{m} \sum_{i=1}^m \sum_{c=1}^C p(y_c | x_i) \log \frac{p(y_c | x_i)}{p(y_c)}\right). \quad (65)$$

It aims at estimating (i) if the samples look realistic i.e., $p(y|x)$ should have low entropy, and (ii) if the samples are diverse (from different ImageNet classes) i.e., $p(y)$ should have high entropy. As these are combined using the Kullback–Leibler divergence, the higher the score is, the better the performance. Note that the range of IS scores at convergence varies across datasets, as the Inception network is pretrained on the ImageNet classes. For example, we obtain low IS values on the SVHN dataset as a large fraction of classes are numbers, which typically do not appear in the ImageNet dataset. Since **MNIST** has greyscale images, we used a classifier trained on this dataset and used $m = 5000$. For the rest of the datasets, we used the original implementation¹⁰ of IS in TensorFlow, and $m = 50000$.

F.1.2 Fréchet Inception Distance

Contrary to IS, FID aims at comparing the synthetic samples $x \sim p_g$ with those of the training dataset $x \sim p_d$ in a feature space. The samples are embedded using the first several layers of the Inception network. Assuming p_g and p_d are multivariate normal distributions, it then estimates the means \mathbf{m}_g and \mathbf{m}_d and covariances C_g and C_d , respectively for p_g and p_d in that feature space. Finally, FID is computed as:

$$\mathbb{D}_{\text{FID}}(p_d, p_g) \approx d^2((\mathbf{m}_d, C_d), (\mathbf{m}_g, C_g)) = \|\mathbf{m}_d - \mathbf{m}_g\|_2^2 + \text{Tr}(C_d + C_g - 2(C_d C_g)^{\frac{1}{2}}), \quad (66)$$

¹⁰<https://github.com/openai/improved-gan/>

where d^2 denotes the Fréchet Distance. Note that as this metric is a distance, the lower it is, the better the performance. We used the original implementation of FID¹¹ in Tensorflow, along with the provided statistics of the datasets.

F.1.3 Second Moment Estimate

To evaluate SVRE effectively, we used the **second moment estimate** (SME, uncentered variance, see § D.1) of the gradient estimate throughout the iterations $t = 1 \dots T$ per parameter, computed as: $v_t = \gamma v_{t-1} + (1-\gamma)g_t^2$, where g_t denotes the gradient estimate for the parameter and iteration t , and $\gamma = 0.9$. For SVRE, g_t is d_φ and d_θ (see Eq. 63 and 64) for G and D , respectively. We initialize $g_0 = 0$ and we use bias-corrected estimates: $\hat{v} = \frac{v_t}{1-\gamma^t}$. As the second moment estimate is computed per each parameter of the model, we depict the average of these values for the parameters of G and D separately.

In this work, as we aim at assessing if SVRE *effectively* reduces the variance of the gradient updates, we use SME in our analysis as it is computationally inexpensive and fast to compute.

F.1.4 Entropy & Total Variation on MNIST

For the experiments on **MNIST** illustrated in Fig. 5a & 3b in § 4, we plot in § G the **entropy** (E) of the generated samples’ class distribution, as well as the **total variation** (TV) between the class distribution of the generated samples and a uniform one (both computed using a pretrained network that classifies its 10 classes).

F.2 Architectures & Hyperparameters

Description of the architectures. We describe the models we used in the empirical evaluation of SVRE by listing the layers they consist of, as adopted in GAN works, e.g. (Miyato et al., 2018). With “conv.” we denote a convolutional layer and “transposed conv” a transposed convolution layer (Radford et al., 2016). The models use Batch Normalization (Ioffe and Szegedy, 2015) and Spectral Normalization layers (Miyato et al., 2018).

F.2.1 Architectures for experiments on MNIST

For experiments on the **MNIST** dataset, we used the DCGAN architectures (Radford et al., 2016), listed in Table 3, and the parameters of the models are initialized using PyTorch default initialization. We used mini-batch sizes of 50 samples, whereas for full dataset passes we used mini-batches of 500 samples as this reduces the wall-clock time for its computation. For experiments on this dataset, we used the *non saturating* GAN loss as proposed (Goodfellow et al., 2014):

$$\mathcal{L}_D = \mathbb{E}_{x \sim p_d} \log(D(x)) + \mathbb{E}_{z \sim p_z} \log(D(G(z))) \quad (67)$$

$$\mathcal{L}_G = \mathbb{E}_{z \sim p_z} \log(D(G(z))), \quad (68)$$

where p_d and p_z denote the data and the latent distributions (the latter to be predefined).

For both the baseline and the SVRE variants we tried the following step sizes $\eta = [1 \times 10^{-2}, 1 \times 10^{-3}, 1 \times 10^{-4}]$. We observe that SVRE can be used with larger step sizes. In Table 9, we used $\eta = 1 \times 10^{-4}$ and $\eta = 1 \times 10^{-2}$ for SE-A and SVRE(-VRAd), respectively.

F.2.2 Choice of architectures on real-world datasets

We replicate the experimental setup described for **CIFAR-10** and **SVHN** in (Miyato et al., 2018), described also below in § F.2.4. We observe that this experimental setup is highly sensitive to the choice of the

¹¹<https://github.com/bioinf-jku/TTUR>

Generator	Discriminator
<i>Input: $z \in \mathbb{R}^{128} \sim \mathcal{N}(0, I)$</i> transposed conv. (ker: 3×3 , $128 \rightarrow 512$; stride: 1) Batch Normalization ReLU transposed conv. (ker: 4×4 , $512 \rightarrow 256$, stride: 2) Batch Normalization ReLU transposed conv. (ker: 4×4 , $256 \rightarrow 128$, stride: 2) Batch Normalization ReLU transposed conv. (ker: 4×4 , $128 \rightarrow 1$, stride: 2, pad: 1) $\text{Tanh}(\cdot)$	<i>Input: $x \in \mathbb{R}^{1 \times 28 \times 28}$</i> conv. (ker: 4×4 , $1 \rightarrow 64$; stride: 2; pad: 1) LeakyReLU (negative slope: 0.2) conv. (ker: 4×4 , $64 \rightarrow 128$; stride: 2; pad: 1) Batch Normalization LeakyReLU (negative slope: 0.2) conv. (ker: 4×4 , $128 \rightarrow 256$; stride: 2; pad: 1) Batch Normalization LeakyReLU (negative slope: 0.2) conv. (ker: 3×3 , $256 \rightarrow 1$; stride: 1) $\text{Sigmoid}(\cdot)$

Table 3: DCGAN architectures (Radford et al., 2016) used for experiments on **MNIST**. We use *ker* and *pad* to denote *kernel* and *padding* for the (transposed) convolution layers, respectively. With $h \times w$ we denote the kernel size. With $c_{in} \rightarrow y_{out}$ we denote the number of channels of the input and output, for (transposed) convolution layers.

hyperparameters (see our results in § G.3), making it more difficult to compare the optimization methods for a fixed hyperparameter choice. In particular, apart from the different combinations of learning rates for G and D , for the baseline this also included experimenting with: β_1 (see (56)), a multiplicative factor of exponential learning rate decay scheduling γ , as well as different ratio of updating G and D per iteration. These observations, combined with that we had limited computational resources, motivated us to use shallower architectures, which we describe below in § F.2.3, and which use an inductive bias of so-called Self-Attention layers (Zhang et al., 2018). As a reference, our SAGAN and ResNet architectures for **CIFAR-10** have approximately 35 and 85 layers, respectively—in total for G and D , including the non linearity and the normalization layers. For clarity, although the deeper and the shallower architectures differ as they are based on ResNet and SAGAN, we refer these as *deep* (see § F.2.3) and *shallow* (see § F.2.4), respectively.

F.2.3 Shallower SAGAN architectures

We used the SAGAN architectures (Zhang et al., 2018), as the techniques of self-attention introduced in SAGAN were used to obtain the state-of-art GAN results on ImageNet (Brock et al., 2019). In summary, these architectures: (i) allow for attention-driven, long-range dependency modeling, (ii) use spectral normalization (Miyato et al., 2018) on both G and D (efficiently computed with the *power iteration* method); and (iii) use different learning rates for G and D , as advocated in (Heusel et al., 2017). The foremost is obtained by combining weights, or alternatively *attention vectors*, with the convolutions across layers, so as to allow modeling textures that are consistent globally—for the generator, or enforcing geometric constraints on the global image structure—for the discriminator.

We used the architectures listed in Table 5 for **CIFAR-10** and **SVHN** datasets, and the architectures described in Table 6 for the experiments on **ImageNet**. The models’ parameters are initialized using the default initialization of PyTorch.

For experiments with SAGAN, we used the hinge version of the adversarial non-saturating loss (Lim and Ye, 2017; Zhang et al., 2018):

$$\mathcal{L}_D = \mathbb{E}_{x \sim p_d} \max(0, 1 - D(x)) + \mathbb{E}_{z \sim p_z} \max(0, 1 + D(G(z))) \quad (69)$$

$$\mathcal{L}_G = -\mathbb{E}_{z \sim p_z} D(G(z)), \quad (70)$$

Self-Attention Block (d – input depth)		
<i>Input:</i> $t \in \mathbb{R}^{d \times H \times W}$		
<i>i:</i> conv. (ker: 1×1 , $d \rightarrow \lfloor d/8 \rfloor$)	<i>ii:</i> conv. (ker: 1×1 , $d \rightarrow \lfloor d/8 \rfloor$)	<i>iii:</i> conv. (ker: 1×1 , $d \rightarrow d$)
<i>iv:</i> softmax($(i) \otimes (ii)$)		
<i>Output:</i> $\gamma((iv) \otimes (iii)) + t$		

Table 4: Layers of the self-attention block used in the SAGAN architectures (see Tables 5 and 6), where \otimes denotes matrix multiplication and γ is a scale parameter initialized with 0. The columns emphasize that the execution is in parallel, more precisely, that the block input t is input to the convolutional layers (i)–(iii). The shown row ordering corresponds to consecutive layers’ order, e.g. softmax is done on the product of the outputs of the (i) and (ii) convolutional layers. The 1×1 convolutional layers have stride of 1. For complete details see Zhang et al. (2018).

Generator		Discriminator	
<i>Input:</i> $z \in \mathbb{R}^{128} \sim \mathcal{N}(0, I)$		<i>Input:</i> $x \in \mathbb{R}^{3 \times 32 \times 32}$	
transposed conv. (ker: 4×4 , $128 \rightarrow 256$; stride: 1)		conv. (ker: 4×4 , $3 \rightarrow 64$; stride: 2; pad: 1)	
Spectral Normalization		Spectral Normalization	
Batch Normalization		LeakyReLU (negative slope: 0.1)	
ReLU		conv. (ker: 4×4 , $64 \rightarrow 128$; stride: 2; pad: 1)	
transposed conv. (ker: 4×4 , $256 \rightarrow 128$, stride: 2, pad: 1)		Spectral Normalization	
Spectral Normalization		LeakyReLU (negative slope: 0.1)	
Batch Normalization		conv. (ker: 4×4 , $128 \rightarrow 256$; stride: 2; pad: 1)	
ReLU		Spectral Normalization	
Self-Attention Block (128)		LeakyReLU (negative slope: 0.1)	
transposed conv. (ker: 4×4 , $128 \rightarrow 64$, stride: 2, pad: 1)		conv. (ker: 4×4 , $256 \rightarrow 1$; stride: 1)	
Spectral Normalization		Self-Attention Block (256)	
Batch Normalization			
ReLU			
Self-Attention Block (64)			
transposed conv. (ker: 4×4 , $64 \rightarrow 3$, stride: 2, pad: 1)			
Tanh(\cdot)			

Table 5: *Shallow* SAGAN architectures for experiments on **SVHN** and **CIFAR-10**, for the Generator (left) and the Discriminator (right). The self-attention block is described in Table 4. We use the default PyTorch hyperparameters for the Batch Normalization layer.

where consistent with the notation above, p_d and p_z denote the data and the latent distributions.

For the SE-A baseline we obtained best performances when $\eta_G = 1 \times 10^{-4}$ and $\eta_D = 4 \times 10^{-4}$, for G and D, respectively. Similarly as noted for **MNIST**, using SVRE allows for using larger order of the step size on the rest of the datasets, whereas SE-A with increased step size ($\eta_G = 1 \times 10^{-3}$ and $\eta_D = 4 \times 10^{-3}$ failed to converge. In Table 2, $\eta_G = 1 \times 10^{-3}$, $\eta_D = 4 \times 10^{-3}$, and $\eta_G = 5 \times 10^{-3}$, $\eta_D = 8 \times 10^{-3}$, $\beta_1 = 0.3$ for SVRE and SVRE-VRad, respectively. We did not use momentum for the vanilla SVRE experiments.

Generator		Discriminator	
<i>Input: $z \in \mathbb{R}^{128} \sim \mathcal{N}(0, I)$</i>		<i>Input: $x \in \mathbb{R}^{3 \times 64 \times 64}$</i>	
transposed conv. (ker: 4×4 , $128 \rightarrow 512$; stride: 1)		conv. (ker: 4×4 , $3 \rightarrow 64$; stride: 2; pad: 1)	
Spectral Normalization		Spectral Normalization	
Batch Normalization		LeakyReLU (negative slope: 0.1)	
ReLU		conv. (ker: 4×4 , $64 \rightarrow 128$; stride: 2; pad: 1)	
transposed conv. (ker: 4×4 , $512 \rightarrow 256$; stride: 2; pad: 1)		Spectral Normalization	
Spectral Normalization		LeakyReLU (negative slope: 0.1)	
Batch Normalization		conv. (ker: 4×4 , $128 \rightarrow 256$; stride: 2; pad: 1)	
ReLU		Spectral Normalization	
transposed conv. (ker: 4×4 , $256 \rightarrow 128$; stride: 2; pad: 1)		LeakyReLU (negative slope: 0.1)	
Spectral Normalization		Self-Attention Block (256)	
Batch Normalization		conv. (ker: 4×4 , $256 \rightarrow 512$; stride: 2; pad: 1)	
ReLU		Spectral Normalization	
Self-Attention Block (128)		LeakyReLU (negative slope: 0.1)	
transposed conv. (ker: 4×4 , $128 \rightarrow 64$; stride: 2; pad: 1)		Self-Attention Block (512)	
Spectral Normalization		conv. (ker: 4×4 , $512 \rightarrow 1$; stride: 1)	
Batch Normalization			
ReLU			
Self-Attention Block (64)			
transposed conv. (ker: 4×4 , $64 \rightarrow 3$; stride: 2; pad: 1)			
$Tanh(\cdot)$			

Table 6: *Shallow* SAGAN architectures for experiments on **ImageNet**, for the Generator (left) and the Discriminator (right). The self-attention block is described in Table 4. Relative to the architectures used for **SVHN** and **CIFAR-10** (see Table 5), the generator has one additional “common” block (conv.–norm.–ReLU), whereas the discriminator has additional “common” block as well as self-attention block (both of more parameters).

F.2.4 Deeper ResNet architectures

We experimented with ResNet (He et al., 2015) architectures on **CIFAR-10** and **SVHN**, using the architectures listed in Table 8, that replicate the setup described in (Miyato et al., 2018) on **CIFAR-10**. For experiments with ResNet, we used the hinge version of the adversarial non-saturating loss, Eq. 69 and 70. For this architectures, we refer the reader to § G.3 for details on the hyperparameters, where we list the hyperparameters along with the obtained results.

G-ResBlock	D-ResBlock (ℓ-th block)
<i>Bypass:</i>	<i>Bypass:</i>
Upsample($\times 2$)	[AvgPool (ker: 2×2)], if $\ell = 1$
<i>Feedforward:</i>	conv. (ker: 1×1 , $3_{\ell=1}/128_{\ell\neq 1} \rightarrow 128$; stride: 1)
Batch Normalization	Spectral Normalization
ReLU	[AvgPool (ker: 2×2 , stride:2)], if $\ell \neq 1$
Upsample($\times 2$)	<i>Feedforward:</i>
conv. (ker: 3×3 , $256 \rightarrow 256$; stride: 1; pad: 1)	[ReLU], if $\ell \neq 1$
Batch Normalization	conv. (ker: 3×3 , $3_{\ell=1}/128_{\ell\neq 1} \rightarrow 128$; stride: 1; pad: 1)
ReLU	Spectral Normalization
conv. (ker: 3×3 , $256 \rightarrow 256$; stride: 1; pad: 1)	ReLU
	conv. (ker: 3×3 , $128 \rightarrow 128$; stride: 1; pad: 1)
	Spectral Normalization
	AvgPool (ker: 2×2)

Table 7: ResNet blocks used for the ResNet architectures (see Table 8), for the Generator (left) and the Discriminator (right). Each ResNet block contains skip connection (bypass), and a sequence of convolutional layers, normalization, and the ReLU non-linearity. The skip connection of the ResNet blocks for the Generator (left) upsamples the input using a factor of 2 (we use the default PyTorch upsampling algorithm—nearest neighbor), whose output is then added to the one obtained from the ResNet block listed above. For clarity we list the layers sequentially, however, note that the bypass layers operate in parallel with the layers denoted as “feedforward” (He et al., 2015). The ResNet block for the Discriminator (right) differs if it is the first block in the network (following the input to the Discriminator), $\ell = 1$, or a subsequent one, $\ell > 1$, so as to avoid performing the ReLU non-linearity immediate on the input.

Generator	Discriminator
<i>Input: $z \in \mathbb{R}^{128} \sim \mathcal{N}(0, I)$</i>	<i>Input: $x \in \mathbb{R}^{3\times 32\times 32}$</i>
Linear($128 \rightarrow 4096$)	D-ResBlock
G-ResBlock	D-ResBlock
G-ResBlock	D-ResBlock
G-ResBlock	D-ResBlock
Batch Normalization	ReLU
ReLU	AvgPool (ker: 8×8)
conv. (ker: 3×3 , $256 \rightarrow 3$; stride: 1; pad:1)	Linear($128 \rightarrow 1$)
$Tanh(\cdot)$	Spectral Normalization

Table 8: Deep ResNet architectures used for experiments on **SVHN** and **CIFAR-10**, where G-ResBlock and D-ResBlock for the Generator (left) and the Discriminator (right), respectively, are described in Table 7. The models’ parameters are initialized using the Xavier initialization (Glorot and Bengio, 2010).

	IS			FID		
	SE-A	SVRE	SVRE-VRAd	SE-A	SVRE	SVRE-VRAd
MNIST	8.62	8.58	8.56	0.17	0.15	0.18
CIFAR-10	6.61	6.50	6.67	37.20	39.20	38.88
SVHN	2.83	3.01	3.04	39.95	24.01	19.40
ImageNet	7.22	8.08	7.50	89.40	75.60	81.24

Table 9: Best obtained IS and FID scores for the different optimization methods, using *shallow* architectures, for a fixed number of iterations (see § F). The architectures for each dataset are described in: **MNIST**–Table 3, **SVHN** and **CIFAR-10**–Table 5, and **ImageNet**–Table 6. The standard deviation of the Inception scores is around 0.1 and is omitted. Although the IS metric gives relatively close values on **SVHN** due to the dataset properties (see § F.1), we include it for completeness.

G Additional Experiments

G.1 Results on MNIST

The results in Table 2 on **MNIST** are obtained using 5 runs with different seeds, and the shown performances are the averaged values. Each experiment was run for 100K iterations. The corresponding scores with the standard deviations are as follows: (i) IS: $8.62 \pm .02$, $8.58 \pm .08$, $8.56 \pm .11$; (ii) FID: $0.17 \pm .03$, $0.15 \pm .01$, $0.18 \pm .02$; for SE-A, SVRE, and SVRE-VRAd, respectively. On this dataset, we obtain similar final performances if run for many iterations, however SVRE converges faster (see Fig. 3). Fig. 5 illustrates additional metrics of the experiments shown in Fig. 3.

G.2 Results with shallow architectures

Fig. 6 depicts the results on **ImageNet** using the *shallow* architectures described in Table 6, § F.2.3. Table 9 summarizes the results obtained on **SVHN**, **CIFAR-10** and **ImageNet** with these architectures. Fig. 7 depicts the SME metric (see § F.1.3) for the SE-A baseline and SVRE shown in Fig. 3c, on **SVHN**.

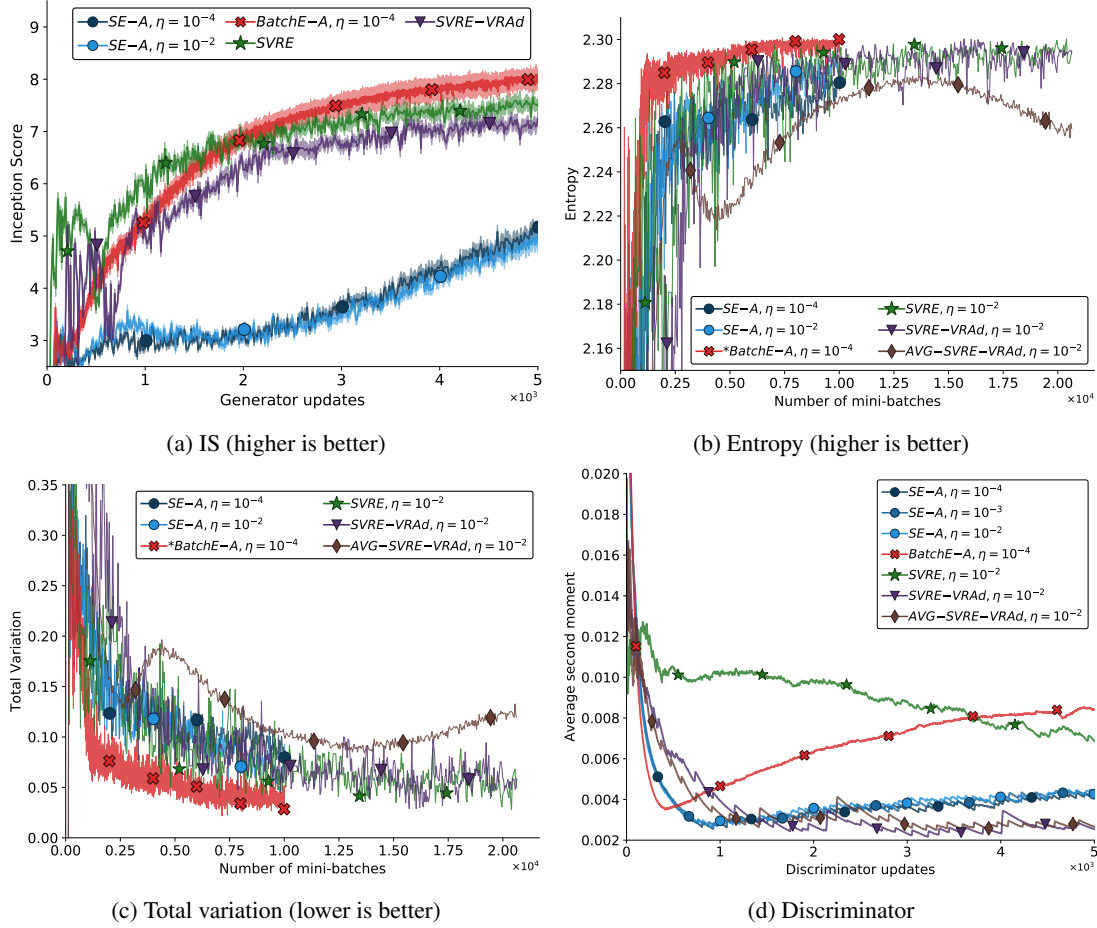


Figure 5: Stochastic, full-batch and variance reduced versions of the extragradient method ran on **MNIST**, see § 4.1. **BatchE-A* emphasizes that this method is **not** scaled with the number of passes (x-axis). The input space is $1 \times 28 \times 28$, see § F.2 for details on the implementation.




Figure 6: Comparison between *SVRE* and the *SE-A* baseline on **Imagenet**, using the *shallow* architectures described in Table 6. See § F.1 for details on the used IS and FID metrics.




Figure 7: Average second moment estimate (see § F.1.3) on **SVHN** for the Generator (left) and the Discriminator (right), using the *shallow* architectures described in Table 5. The corresponding FID scores for these experiments are shown in Fig. 3c.

G.3 Results with deeper architectures

We observe that GAN training is more challenging when using *deeper* architectures and some empirical observations differ in the two settings. For example, our stochastic baseline is drastically more unstable and often does not start to converge, whereas SVRE is notably *stable*, but slower compared to when using shallower architectures. In this section, all our discussions focus on *deep* architectures (see § F.2.4).

Stability: convergence of the GAN training. For our stochastic baselines, *irrespective whether we use the extragradient or gradient method*, we observe that the convergence is notably more *unstable* (see Fig. 8) when using the *deep* architectures described in § F.2.4. More precisely, either the training fails to converge or it diverges at later iterations. When updating G and D equal number of times *i.e.* using 1 : 1 update ratio, using SE-A on **CIFAR-10** we obtained best FID score of 24.91 using $\eta_G = 2 \times 10^{-4}$, $\eta_D = 4 \times 10^{-4}$, $\beta_1 = 0$, while experimenting with several combinations of η_G, η_D, β_1 . Using exponential learning rate decay with a multiplicative factor of 0.99, improved the best FID score to 20.70, obtained for the experiment with $\eta_G = 2 \times 10^{-4}$, $\eta_D = 2 \times 10^{-4}$, $\beta_1 = 0$. Finally, using 1 : 5 update ratio, with $\eta_G = 2 \times 10^{-4}$, $\eta_D = 2 \times 10^{-4}$, $\beta_1 = 0$ provided best FID of 18.65 for the baseline. Figures 8a and 8b depict the hyperparameter sensitivity of SE-A and SG-A, respectively. The latter denotes the alternating GAN training with Adam, that is most commonly used for GAN training.




Figure 8: FID scores (lower is better) with different hyperparameters for the SE-A baseline on **CIFAR10** (left) and the SG-A baseline on **SVHN** (right), using the *deep* architectures described in Table 8, § F.2.4. SG-A denotes the standard stochastic *alternating* GAN training, with the Adam optimization method. Where omitted, $\beta_1 = 0$, see (56) where this hyperparameter is defined. With r we denote the update ratio of generator versus discriminator: in particular 1 : 5 denotes that D is updated 5 times for each update of G . γ denotes a multiplicative factor of exponential learning rate decay scheduling. In Fig. 8b, $\gamma = 0.99$ for all the experiments. We observed in all our experiments that training diverged in later iterations for the stochastic baseline, when using *deep* architectures.

We observe that SVRE is more stable in terms of hyperparameter selection, as it always starts to converge and *does not diverge* at later iterations. Relative to experiments with shallower architectures, we observe that with deeper architectures SVRE takes longer to converge than its baseline for this architecture. With constant step size of $\eta_G = 1 \times 10^{-3}$, $\eta_D = 4 \times 10^{-3}$ we obtain FID score of 23.56 on **CIFAR-10**. Note that this result outperforms the baseline when using no additional tricks (which themselves require additional hyperparameter tuning). Fig. 9 depicts the FID scores obtained when training with SVRE on the **SVHN** dataset, for two




Figure 9: Obtained FID (lower is better) scores for SVRE, using the *deep* architectures (see § F.2.4) on **SVHN**. With s we denote the fixed random seed. The update ratio for all the experiments is 1 : 1. We illustrate our results on the same plot (besides the reduced clarity) so as to summarize our observation that, contrary to the SE-A baseline for these architectures, SVRE *always converges*, and does not diverge.




Figure 10: Obtained FID (lower is better) scores for WS-SVRE, using the *deep* architectures (see § F.2.4) on **CIFAR-10**, where the seed is fixed to 1 for all the experiments. With r we denote the update ratio of generator versus discriminator: in particular 1 : 5 denotes that D is updated 5 times for each update of G . We start from the best obtained FID score for the stochastic baseline, i.e. FID of 18.65 (see Table 2)—shown with dashed line, and we continue to train with SVRE.

different hyperparameter settings, using four different seeds for each. From this set of experiments, we observe that contrary to the baseline that either did not converge or diverged in all our experiments, SVRE always converges. However, we observe different performances for different seeds. This suggests that more exhaustive empirical hyperparameter search that aims to find an empirical setup that works *best* for SVRE or further combining SVRE with adaptive step size techniques are both promising research directions (see our discussion below). Fig. 10 depicts our WS–SVRE experiment, where we start from a stored checkpoint for which we obtained best FID score for the SE–A baseline, and we continue the training with SVRE. It is interesting that besides that the baseline diverged after the stored checkpoint, SVRE further reduced the FID score. Moreover, we observe that using different update ratios does not impact much the performance, what on the other hand was necessary to make the baseline algorithm converge.

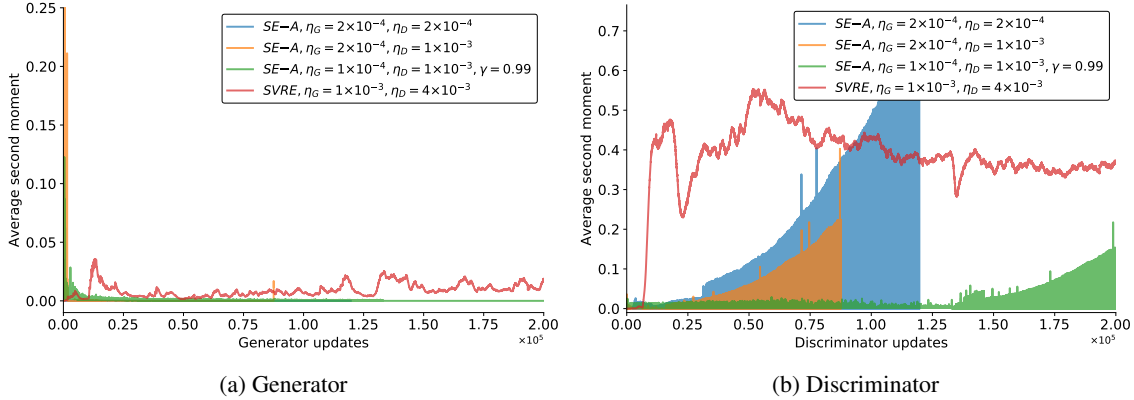


Figure 11: Average second moment estimate (SME, see § F.1.3) on **CIFAR-10** for the Generator (left) and the Discriminator (right), using the *deep* architectures described in Table 8. The obtained FID scores for these experiments are shown in Fig. 8a, where we omit some of the experiments for clarity. All of the baseline SE–A experiments diverge at some point, what correlates with the iterations at which large oscillations of SME appear for the Discriminator. Note that the SE–A experiments were stopped after the algorithm diverges, hence the plotted SME is up to a particular iteration for two of the experiments (shown in blue and orange). The SE–A experiment with $\gamma = 0.99$ diverged at later iteration relative to the experiments without learning rate decay, and has lower SME.

Second moment estimate (SME). Fig. 11 depicts the second moment estimate (see § F.1.3) for the experiments with *deep* architectures. We observe that: (i) the estimated SME quantity is more bounded and changes more smoothly for SVRE (as we do not observe large oscillations of it as it is the case for SE–A); as well as that (ii) divergence of the SE–A baseline *correlates* with large oscillations of SME, in this case, observed for the Discriminator. Regarding the latter, there exist larger in magnitude oscillations of SME (note that the exponential moving average hyperparameter for computing SME is $\gamma = 0.9$, see § F.1.3).

Conclusion & future directions. In summary, we observe the following most important advantages of SVRE when using *deep* architectures: (i) consistency of convergence, and improved stability; as well as (ii) reduced number of hyperparameters. Apart from the practical benefit for applications, the former could allow for a more fair comparison of GAN variants. The latter refers to the fact that SVRE omits the tuning of the sensitive (for the stochastic baseline) β_1 hyperparameter (see (56)), as well as r and γ —as training converges for SVRE without using different update ratio and step size schedule, respectively. It is important to note that the stochastic baseline does not converge when using constant step size (i.e. when *SGD* is used instead of *Adam*). In our experiments we compared SVRE that uses constant step size, with *Adam*, making

the comparison unfair toward SVRE. Hence, our results indicate that SVRE can be further combined with adaptive step size schemes, so as to obtain both stable GAN performances and fast convergence when using these architectures. Nonetheless, the fact that the baseline either does not start to converge or it diverges later makes SVRE and WS-SVRE a promising approach for practitioners using these *deep* architectures, whereas, for *shallower* ones, SVRE speeds up the convergence and often provides better *final* performances.