

ON LYAPUNOV FUNCTIONS AND PARTICLE METHODS FOR REGULARIZED MINIMAX PROBLEMS

LEXING YING

ABSTRACT. We study the two-player zero-sum game with mixed strategies. For a class of commonly-used regularizers and a class of metrics, we show the existence of a Lyapunov function of the gradient ascent descent dynamics. We also propose for a new particle method for a specific combination of regularizers and metrics.

1. INTRODUCTION

This note is concerned with two-player zero-sum game with mixed strategies. Let Ω_1 and Ω_2 be two compact sets of strategies and $K(x_1, x_2)$ for $x_1 \in \Omega_1$ $x_2 \in \Omega_2$ be the payoff function. $\mathcal{P}(\Omega_1)$ and $\mathcal{P}(\Omega_2)$ denote the spaces of probability densities over Ω_1 and Ω_2 , respectively. When $K(x_1, x_2)$ is continuous, the two-layer zero-sum game with mixed strategies and payoff

$$p_1^\top K p_2 \equiv \iint_{\Omega_1 \times \Omega_2} p_1(x_1) K(x_1, x_2) p_2(x_2) dx_1 dx_2, \quad p_1 \in \mathcal{P}(\Omega_1), p_2 \in \mathcal{P}(\Omega_2)$$

has a unique Nash equilibrium [4] given by

$$\min_{p_1 \in \mathcal{P}(\Omega_1)} \max_{p_2 \in \mathcal{P}(\Omega_2)} p_1^\top K p_2 = \max_{p_2 \in \mathcal{P}(\Omega_2)} \min_{p_1 \in \mathcal{P}(\Omega_1)} p_1^\top K p_2.$$

Due to stability, it is often useful to consider a regularized version

$$\min_{p_1} \max_{p_2} H_1(p_1) + p_1^\top K p_2 - H_2(p_2),$$

where $H_1(p_1)$ and $H_2(p_2)$ are the regularizers applied to $p_1(x_1)$ and $p_2(x_2)$ or the more general form

$$(1) \quad \min_{p_1} \max_{p_2} E(p_1, p_2) \equiv H_1(p_1) + e_1^\top p_1 + p_1^\top K p_2 - H_2(p_2) - e_2^\top p_2$$

with the extra linear terms $e_1^\top p_1 \equiv \int_{\Omega_1} e_1(x_1) p_1(x_1) dx_1$ and $e_2^\top p_2 \equiv \int_{\Omega_2} e_2(x_2) p_2(x_2) dx_2$.

This note studies the gradient ascent descent (GAD) dynamics for solving (1). The main contributions are listed as follows.

- For a general class of regularizers and for a general class of metrics, we show the existence of a Lyapunov function for the GAD of (1).
- For a specific combination of the regularizers and the metrics, we propose a new particle method for solving (1).

Key words and phrases. Path divergence, Wasserstein gradient descent ascent, minimax problems.

The work of L.Y. is partially supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Scientific Discovery through Advanced Computing (SciDAC) program and also by the National Science Foundation under award DMS-2011699.

Related work. The two-player zero-sum game has received a lot of attention in machine learning since the introduction of the generative adversarial networks (GANs) [5]. Most of the works inspired by GANs have focused on the pure strategy case, either for the convex-concave games [9, 11, 12] or for the local equilibria [1, 6, 8, 10]. In the area of mixed strategies, this note is mostly inspired by the recent works in [2, 7]. [7] studied the mirror ascent descent dynamics of the non-regularized problem and proposed an implementation based on running a Langevin dynamics at each time step. [2] studied the dynamics under the Wasserstein-Fisher-Rao metric for the non-regularized problem and proved finite-time error bounds in the weak transport regime. Compared with these two papers, the current note focuses on the dynamics with non-transport metrics for the regularized problems and also proposes a new particle method.

Contents. The rest of the note is organized as follows. Section 2 describes the general setup and proves the existence of the Lyapunov function for the gradient ascent descent dynamics. Section 3 studies a few special cases, proposes a new particle method, and provides on some extensions. Section 4 concludes with some discussion for future directions.

2. LYAPUNOV FUNCTION

We consider the regularizers of the form

$$H_1(p_1) = \int_{\Omega_1} h_1(p_1(x_1))dx_1, \quad H_2(p_2) = \int_{\Omega_2} h_2(p_2(x_2))dx_2,$$

where $h_1(\cdot)$ and $h_2(\cdot)$ are strictly convex functions defined on the positive real axis. With these regularizers, the objective function $E(p_1, p_2)$ in (1) takes the following form

$$(2) \quad E(p_1, p_2) = \int h_1(p_1(x_1))dx_1 + e_1^\top p_1 + p_1^\top K p_2 - \int h_2(p_2(x_2))dx_2 - e_2^\top p_2.$$

The functional derivatives of $E(p_1, p_2)$ in p_1 and p_2 are respectively

$$\delta_{p_1} E(p_1, p_2) = +h_1'(p_1) + e_1 + K p_2, \quad \delta_{p_2} E(p_1, p_2) = -h_2'(p_2) - e_2 + K^\top p_1.$$

Let us introduce the following metric functionals for $p_1 \in \mathcal{P}(\Omega_1)$ and $p_2 \in \mathcal{P}(\Omega_2)$

$$M_1(p_1) = \int_{\Omega_1} m_1(p_1(x_1))dx_1, \quad M_2(p_2) = \int_{\Omega_2} m_2(p_2(x_2))dx_2,$$

where $m_1(\cdot)$ and $m_2(\cdot)$ are strictly convex functions over the positive real axis. The Hessians of these metric functionals

$$\delta_{p_1 p_1} M_1(p_1) = \text{diag}(m_1''(p_1)), \quad \delta_{p_2 p_2} M_2(p_2) = \text{diag}(m_2''(p_2)),$$

introduce non-Euclidean metrics on the spaces $\mathcal{P}(\Omega_1)$ and $\mathcal{P}(\Omega_2)$, respectively.

The gradient ascent descent of $E(p_1, p_2)$ under these metrics is given by

$$(3) \quad \begin{aligned} \partial_t p_1 &= - (m_1''(p_1))^{-1} (h_1'(p_1) + e_1 + K p_2 + \text{cst}), \\ \partial_t p_2 &= - (m_2''(p_2))^{-1} (h_2'(p_2) + e_2 - K^\top p_1 + \text{cst}) \end{aligned}$$

where the constant cst is introduced to ensure that p_1 and p_2 both remain to be probability distributions, i.e., $\int_{\Omega_1} p_1(x_1)dx_1 = \int_{\Omega_2} p_2(x_2)dx_2 = 1$.

The rest of this section is to show that (3) has a Lyapunov function. Since Ω_1 and Ω_2 are compact and $K(x_1, x_2)$ is continuous, (1) has a unique solution (see [4]), which shall be denoted by (p_1^*, p_2^*) in what follows. The first order optimality condition of (2) states that

$$(4) \quad \begin{aligned} h_1'(p_1^*) + e_1 + Kp_2^* &= \text{cst}, \\ h_2'(p_2^*) + e_2 + K^\top p_1^* &= \text{cst}. \end{aligned}$$

The Bregman divergences of $M_1(p_1)$ and $M_2(p_2)$ based at p_1^* and p_2^* are given by

$$(5) \quad \begin{aligned} D_{M_1}(p_1^*, p_1) &= M_1(p_1^*) - M_1(p_1) - \langle \delta_{p_1} M_1(p_1), p_1^* - p_1 \rangle \\ D_{M_2}(p_2^*, p_2) &= M_2(p_2^*) - M_2(p_2) - \langle \delta_{p_2} M_2(p_2), p_2^* - p_2 \rangle. \end{aligned}$$

In what follows, we fix p_1^* and p_2^* and consider them only as functions of p_1 and p_2 . These Bregman divergences are equal to zero if and only if $p_1 = p_1^*$ and $p_2 = p_2^*$, respectively, due to the strict convexity of $m_1(\cdot)$ and $m_2(\cdot)$.

Theorem 1. $L(p_1, p_2) \equiv D_{M_1}(p_1^*, p_1) + D_{M_2}(p_2^*, p_2)$ is a Lyapunov function for the dynamics in (3).

Proof. Subtracting (4) from the right-hand-sides of (3), we obtain

$$(6) \quad \begin{aligned} \partial_t p_1 &= - (m_1''(p_1))^{-1} (h_1'(p_1) - h_1'(p_1^*) + K(p_2 - p_2^*) + \text{cst}), \\ \partial_t p_2 &= - (m_2''(p_2))^{-1} (h_2'(p_2) - h_2'(p_2^*) - K^\top(p_1 - p_1^*) + \text{cst}). \end{aligned}$$

The functional derivatives of $D_{M_1}(p_1^*, p_1)$ and $D_{M_2}(p_2^*, p_2)$ in p_1 and p_2 are, respectively

$$\delta_{p_1} D_{M_1}(p_1^*, p_1) = (p_1 - p_1^*) m_1''(p_1), \quad \delta_{p_2} D_{M_2}(p_2^*, p_2) = (p_2 - p_2^*) m_2''(p_2).$$

The time derivative $d_t L(p_1(t), p_2(t))$ is given by

$$\begin{aligned} & \langle \delta_{p_1} D_{M_1}(p_1^*, p_1), \partial_t p_1 \rangle + \langle \delta_{p_2} D_{M_2}(p_2^*, p_2), \partial_t p_2 \rangle \\ &= - \int (p_1 - p_1^*) m_1''(p_1) (m_1''(p_1))^{-1} (h_1'(p_1) - h_1'(p_1^*) + K(p_2 - p_2^*) + \text{cst}) dx_1 \\ & \quad - \int (p_2 - p_2^*) m_2''(p_2) (m_2''(p_2))^{-1} (h_2'(p_2) - h_2'(p_2^*) - K^\top(p_1 - p_1^*) + \text{cst}) dx_2 \\ &= - \int (p_1 - p_1^*) (h_1'(p_1) - h_1'(p_1^*)) dx_1 - \int (p_2 - p_2^*) (h_2'(p_2) - h_2'(p_2^*)) dx_2. \end{aligned}$$

Since $h_1'(p_1)$ and $h_2'(p_2)$ is strictly monotone, the last quantity is strictly less than zero, except at $p_1 = p_1^*$ and $p_2 = p_2^*$. Therefore, $L(p_1, p_2)$ is a Lyapunov function for the dynamics in (3). \square

3. SPECIAL CASES AND EXTENSIONS

The result in Section 2 holds for rather general functions $h_1(p_1)$, $h_2(p_2)$, $m_1(p_1)$, and $m_2(p_2)$. This section studies a few special cases.

3.1. Regularizer equal to metric functional. In this case, $h_1(p_1) = m_1(p_1)$ and $h_2(p_2) = m_2(p_2)$, which leads to the dynamics

$$(7) \quad \begin{aligned} \partial_t p_1 &= - (h_1''(p_1))^{-1} (h_1'(p_1) + e_1 + K p_2 + \text{cst}), \\ \partial_t p_2 &= - (h_2''(p_2))^{-1} (h_2'(p_2) + e_2 - K^\top p_1 + \text{cst}), \end{aligned}$$

or more conveniently in the vector form

$$(8) \quad \partial_t \begin{pmatrix} p_1 \\ p_2 \end{pmatrix} = - \begin{pmatrix} h_1''(p_1) & \\ & h_2''(p_2) \end{pmatrix}^{-1} \begin{pmatrix} h_1'(p_1) + e_1 + K p_2 + \text{cst} \\ h_2'(p_2) + e_2 - K^\top p_1 + \text{cst} \end{pmatrix}.$$

As the Newton dynamics for solving the stationary condition (4) is

$$\partial_t \begin{pmatrix} p_1 \\ p_2 \end{pmatrix} = - \begin{pmatrix} h_1''(p_1) & K \\ K^\top & h_2''(p_2) \end{pmatrix}^{-1} \begin{pmatrix} h_1'(p_1) + e_1 + K p_2 + \text{cst} \\ h_2'(p_2) + e_2 - K^\top p_1 + \text{cst} \end{pmatrix},$$

we can view (8) as a block diagonal approximation to the Newton dynamics. One key advantage of working with (8) is that it can also be written as

$$\partial_t \begin{pmatrix} h_1'(p_1) \\ h_2'(p_2) \end{pmatrix} = - \begin{pmatrix} h_1'(p_1) + e_1 + K p_2 + \text{cst} \\ h_2'(p_2) + e_2 - K^\top p_1 + \text{cst} \end{pmatrix}.$$

By working directly with $h_1'(p_1)$ and $h_2'(p_2)$, one can discretize with large time steps.

The most important example is $h_1(p) = m_1(p) = p \log p$ and $h_2(p) = m_2(p) = p \log p$. With these choices

$$h_1'(p) = \log p + \text{cst}, \quad h_2'(p) = \log p + \text{cst}, \quad m_1''(p) = 1/p, \quad m_2''(p) = 1/p,$$

and the dynamics (7) becomes

$$(9) \quad \begin{aligned} \partial_t p_1 &= -p_1 (\log p_1 + e_1 + K p_2 + \text{cst}), \\ \partial_t p_2 &= -p_2 (\log p_2 + e_2 - K^\top p_1 + \text{cst}) \end{aligned}$$

or equivalently in terms of $\log p_1$ and $\log p_2$

$$\begin{aligned} \partial_t \log p_1 &= -(\log p_1 + e_1 + K p_2 + \text{cst}), \\ \partial_t \log p_2 &= -(\log p_2 + e_2 - K^\top p_1 + \text{cst}). \end{aligned}$$

An explicit discretization with time step Δt leads to the mirror ascent descent algorithm

$$(10) \quad \begin{aligned} p_1(t + \Delta t) &\propto p_1(t)^{1-\Delta t} \cdot e^{-\Delta t(e_1 + K p_2(t))} \\ p_2(t + \Delta t) &\propto p_2(t)^{1-\Delta t} \cdot e^{-\Delta t(e_2 - K^\top p_1(t))}, \end{aligned}$$

where \propto means proportional to, i.e., a normalization step is required to ensure $\int_{\Omega_1} p_1(x_1) dx_1 = \int_{\Omega_2} p_2(x_2) dx_2 = 1$.

As an approximation to the Newton dynamics, (9) and (10) can offer fast convergence when Ω_1 and Ω_2 can be discretized easily. To illustrate this, Figure 1 gives a simple one-dimensional example with Ω_1 and Ω_2 given by the periodic interval $[0, 1]$. The left plot shows the payoff function $K(x_1, x_2)$ and the middle plot gives the solution pair $(p_1^*(x_1), p_2^*(x_2))$. The domains Ω_1 and Ω_2 are discretized with a uniform grid of 128 points and the time step Δt is taken to be equal to 1. The iteration in (10) converges in about 40 iterations and the right plot displays how the Lyapunov function decays with respect to the iteration count.




FIGURE 1. Mirror ascent descent algorithm. (a) the payoff function $K(x_1, x_2)$. (b) the optimal solution pair $(p_1^*(x_1), p_2^*(x_2))$. (c) The value of the Lyapunov function as a function of the iteration count.

When Ω_1 and Ω_2 cannot be discretized easily (especially when Ω_1 and Ω_2 are high-dimensional), it is often difficult to work with (9). Though there exists particle methods for (9) based on the birth-death process, the fact that no particles are introduced at new locations in the birth-death process severely constrains its applicability. Another issue with this particle method is that it requires density estimation of $p_1(x_1)$ and $p_2(x_2)$ at the particle locations, which can be computationally expensive when the number of particles is large.

3.2. A new particle method. This subsection introduces a new particle method for (3). We choose

$$h_1(p) = \log 1/p, \quad h_2(p) = \log 1/p, \quad m_1(p) = p \log p, \quad m_2(p) = p \log p.$$

This specific choice gives rise to

$$h'_1(p) = -1/p, \quad h'_2(p) = -1/p, \quad m''_1(p) = 1/p, \quad m''_2(p) = 1/p.$$

The dynamics associated with this choice is

$$(11) \quad \begin{aligned} \partial_t p_1 &= -p_1 (-1/p_1 + e_1 + K p_2 + \text{cst}) = -p_1 (e_1 + K p_2 + \text{cst}) + 1, \\ \partial_t p_2 &= -p_2 (-1/p_2 + e_2 - K^\top p_1 + \text{cst}) = -p_2 (e_2 - K^\top p_1 + \text{cst}) + 1. \end{aligned}$$

This dynamics can be implemented with a particle method, where

- the terms proportional to p_1 or p_2 can be realized with a birth-death process,
- the constant 1 terms can be realized by injecting new particle randomly into Ω_1 and Ω_2 .

Compared with the particle method associated with (9), this method introduces particles at new locations and requires no density estimation. The algorithm is detailed in Algorithm 1, where $\{x_{1,i}\}_{i=1,\dots,n}$ are the particles for $p_1(x_1)$ and $\{x_{2,j}\}_{j=1,\dots,n}$ are the ones for $p_2(x_2)$.

3.3. Extension. The discussion in Section 2 can also be extended to the case where the regularizers and metric functionals are f -divergences [13]. Let us consider the general regularizers

$$H_1(p_1) = \int_{\Omega_1} h_1 \left(\frac{p_1(x_1)}{\mu_1(x_1)} \right) \mu_1(x_1) dx_1, \quad H_2(p_2) = \int_{\Omega_2} h_2 \left(\frac{p_2(x_2)}{\mu_2(x_2)} \right) \mu_2(x_2) dx_2,$$

Algorithm 1 Particle algorithm for (11)

Require: Initialize $\{x_{1,i}\}_{i=1,\dots,n}$ from Ω_1 and $\{x_{2,j}\}_{j=1,\dots,n}$ from Ω_2 . Time step Δt .

while until convergence **do**
 for $i = 1, \dots, n$ **do**
 $\alpha_i = (-\frac{1}{n} \sum_j K(x_{1,i}, x_{2,j}) + e_1(x_{1,i}))$.
 end for
 Subtract from each α_i the average $\frac{1}{n} \sum_k \alpha_k$.
 for $j = 1, \dots, n$ **do**
 $\beta_j = (+\frac{1}{n} \sum_i K(x_{1,i}, x_{2,j}) + e_2(x_{2,j}))$.
 end for
 Subtract from each β_j the average $\frac{1}{n} \sum_k \beta_k$.
 for $i = 1, \dots, n$ **do**
 If $\alpha_i < 0$, kill $x_{1,i}$ with probability $1 - \exp(\alpha_i \Delta t)$.
 If $\alpha_i > 0$, duplicate $x_{1,i}$ with probability $1 - \exp(-\alpha_i \Delta t)$.
 end for
 Resample n samples from $\{x_{1,i}\}$ and define them to be $\{x_{1,i}\}$.
 for $j = 1, \dots, n$ **do**
 If $\beta_j < 0$, kill $x_{2,j}$ with probability $1 - \exp(\beta_j \Delta t)$.
 If $\beta_j > 0$, duplicate $x_{2,j}$ with probability $1 - \exp(-\beta_j \Delta t)$.
 end for
 Resample n samples from $\{x_{2,j}\}$ and define them to be $\{x_{2,j}\}$.
 for $i = 1, \dots, n$ **do**
 Keep each $x_{1,i}$ with probability $\exp(-\Delta t)$
 If $x_{1,i}$ is deleted, replace it with a uniform sample from Ω_1 .
 end for
 for $j = 1, \dots, n$ **do**
 Keep each $x_{2,j}$ with probability $\exp(-\Delta t)$
 If $x_{2,j}$ is deleted, replace it with a uniform sample from Ω_2 .
 end for
end while

where $\mu_1(x_1)$ and $\mu_2(x_2)$ are positive reference densities on Ω_1 and Ω_2 , respectively. The objective function is

$$(12) \quad E(p_1, p_2) = \int h_1 \left(\frac{p_1(x_1)}{\mu_1(x_1)} \right) \mu_1(x_1) dx_1 + e_1^\top p_1 + p_1^\top K p_2 - \int h_2 \left(\frac{p_2(x_2)}{\mu_2(x_2)} \right) \mu_2(x_2) dx_2 - e_2^\top p_2.$$

The functional derivatives of $E(p_1, p_2)$ in p_1 and p_2 are

$$\delta_{p_1} E(p_1, p_2) = +h_1'(p_1/\mu_1) + e_1 + K p_2, \quad \delta_{p_2} E(p_1, p_2) = -h_2'(p_2/\mu_2) - e_2 + K^\top p_1.$$

Consider the metric functionals

$$M_1(p_1) = \int_{\Omega_1} m_1 \left(\frac{p_1(x_1)}{\nu_1(x_1)} \right) \nu_1(x_1) dx_1, \quad M_2(p_2) = \int_{\Omega_2} m_2 \left(\frac{p_2(x_2)}{\nu_2(x_2)} \right) \nu_2(x_2) dx_2,$$

where $\nu_1(x_1)$ and $\nu_2(x_2)$ are again positive reference densities on Ω_1 and Ω_2 . Note that $\mu_1(x_1)$ and $\nu_1(x_1)$ can be different and the same applies to $\mu_2(x_2)$ and $\nu_2(x_2)$. The Hessians of these

metric functionals are

$$\delta_{p_1 p_1} M_1(p_1) = \text{diag}(m_1''(p_1/\nu_1)/\nu_1), \quad \delta_{p_2 p_2} M_2(p_2) = \text{diag}(m_2''(p_2/\nu_2)/\nu_2).$$

The gradient ascent descent for $E(p_1, p_2)$ under these metrics is given by

$$(13) \quad \begin{aligned} \partial_t p_1 &= -\nu_1 (m_1''(p_1/\nu_1))^{-1} (h_1'(p_1/\mu_1) + e_1 + K p_2 + \text{cst}), \\ \partial_t p_2 &= -\nu_2 (m_2''(p_2/\nu_2))^{-1} (h_2'(p_2/\mu_2) + e_2 - K^\top p_1 + \text{cst}). \end{aligned}$$

The unique solution of the minimax problem of $E(p_1, p_2)$, denoted by (p_1^*, p_2^*) , satisfies the first order optimality condition

$$(14) \quad \begin{aligned} h_1'(p_1^*/\mu_1) + e_1 + K p_2^* &= \text{cst}, \\ h_2'(p_2^*/\mu_2) + e_2 + K^\top p_1^* &= \text{cst}. \end{aligned}$$

The Bregman divergences of $M_1(p_1)$ and $M_2(p_2)$ with respect to p_1^* and p_2^* are

$$(15) \quad \begin{aligned} D_{M_1}(p_1^*, p_1) &= \int m_1(p_1^*/\nu_1)\nu_1 - m_1(p_1/\nu_1)\nu_1 - (p_1^* - p_1)m_1'(p_1/\nu_1)dx_1, \\ D_{M_2}(p_2^*, p_2) &= \int m_2(p_2^*/\nu_2)\nu_2 - m_2(p_2/\nu_2)\nu_2 - (p_2^* - p_2)m_2'(p_2/\nu_2)dx_2. \end{aligned}$$

The functional derivatives of $D_{M_1}(p_1^*, p_1)$ and $D_{M_2}(p_2^*, p_2)$ are

$$\delta_{p_1} D_{M_1}(p_1^*, p_1) = (p_1 - p_1^*)m_1''(p_1/\nu_1)/\nu_1, \quad \delta_{p_2} D_{M_2}(p_2^*, p_2) = (p_2 - p_2^*)m_2''(p_2/\nu_2)/\nu_2.$$

The following calculation shows that the sum $L(p_1, p_2) = D_{M_1}(p_1^*, p_1) + D_{M_2}(p_2^*, p_2)$ is a Lyapunov function for the dynamics (13):

$$\begin{aligned} d_t L(p_1(t), p_2(t)) &= \langle \delta_{p_1} D_{M_1, p_1^*}(p_1), \partial_t p_1 \rangle + \langle \delta_{p_2} D_{M_2, p_2^*}(p_2), \partial_t p_2 \rangle \\ &= - \int (p_1 - p_1^*)m_1''(p_1/\nu_1)/\nu_1 \cdot \nu_1 (m_1''(p_1/\nu_1))^{-1} (h_1'(p_1/\mu_1) - h_1'(p_1^*/\mu_1) + K(p_2 - p_2^*) + \text{cst}) dx_1 \\ &\quad - \int (p_2 - p_2^*)m_2''(p_2/\nu_2)/\nu_2 \cdot \nu_2 (m_2''(p_2/\nu_2))^{-1} (h_2'(p_2/\mu_2) - h_2'(p_2^*/\mu_2) - K^\top(p_1 - p_1^*) + \text{cst}) dx_2 \\ &= - \int (p_1 - p_1^*) (h_1'(p_1/\mu_1) - h_1'(p_1^*/\mu_1)) dx_1 - \int (p_2 - p_2^*) (h_2'(p_2/\mu_2) - h_2'(p_2^*/\mu_2)) dx_2, \end{aligned}$$

which is less than zero due to the strict convexity of h_1 and h_2 , except at $p_1 = p_1^*$ and $p_2 = p_2^*$.

4. DISCUSSIONS

Though the particle method described in Section 3.2 introduces new particles at random locations, the method is not very efficient for high-dimensional problems since these inserted particles do not move. Another dynamics in the literature is the Wasserstein ascent descent

$$(16) \quad \begin{aligned} \partial_t p_1 &= (\text{div } p_1 \nabla)(\ln p_1 + e_1 + K p_2) = \Delta p_1 + \text{div}(p_1 \nabla(e_1 + K p_2)), \\ \partial_t p_2 &= (\text{div } p_2 \nabla)(\ln p_2 + e_2 - K^\top p_1) = \Delta p_2 + \text{div}(p_2 \nabla(e_2 - K^\top p_1)). \end{aligned}$$

The convergence property of (16) is more subtle [3]. An interesting direction following [2] is whether combining the dynamics introduced in this note with (16) would improve its convergence behavior.

REFERENCES

- [1] Leonard Adolphs, Hadi Daneshmand, Aurelien Lucchi, and Thomas Hofmann, *Local saddle point optimization: A curvature exploitation approach*, The 22nd international conference on artificial intelligence and statistics, 2019, pp. 486–495.
- [2] Carles Domingo-Enrich, Samy Jelassi, Arthur Mensch, Grant Rotskoff, and Joan Bruna, *A mean-field analysis of two-player zero-sum games*, arXiv preprint arXiv:2002.06277 (2020).
- [3] Andreas Eberle, Arnaud Guillin, and Raphael Zimmer, *Quantitative harris-type theorems for diffusions and mckean–ulasov processes*, Transactions of the American Mathematical Society **371** (2019), no. 10, 7135–7173.
- [4] Irving L Glicksberg, *A further generalization of the kakutani fixed point theorem, with application to nash equilibrium points*, Proceedings of the American Mathematical Society **3** (1952), no. 1, 170–174.
- [5] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio, *Generative adversarial nets*, Advances in neural information processing systems **27** (2014).
- [6] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter, *Gans trained by a two time-scale update rule converge to a local nash equilibrium*, Advances in neural information processing systems **30** (2017).
- [7] Ya-Ping Hsieh, Chen Liu, and Volkan Cevher, *Finding mixed nash equilibria of generative adversarial networks*, International conference on machine learning, 2019, pp. 2810–2819.
- [8] Chi Jin, Praneeth Netrapalli, and Michael I Jordan, *Minmax optimization: Stable limit points of gradient descent ascent are locally optimal*, arXiv preprint arXiv:1902.00618 (2019).
- [9] Qihang Lin, Mingrui Liu, Hassan Rafique, and Tianbao Yang, *Solving weakly-convex-weakly-concave saddle-point problems as weakly-monotone variational inequality*, arXiv preprint arXiv:1810.10207 **5** (2018).
- [10] Eric V Mazumdar, Michael I Jordan, and S Shankar Sastry, *On finding local nash equilibria (and only local nash equilibria) in zero-sum games*, arXiv preprint arXiv:1901.00838 (2019).
- [11] Panayotis Mertikopoulos, Bruno Lecouat, Houssam Zenati, Chuan-Sheng Foo, Vijay Chandrasekhar, and Georgios Piliouras, *Optimistic mirror descent in saddle-point problems: Going the extra (gradient) mile*, arXiv preprint arXiv:1807.02629 (2018).
- [12] Maher Nouiehed, Maziar Sanjabi, Tianjian Huang, Jason D Lee, and Meisam Razaviyayn, *Solving a class of non-convex min-max games using iterative first order methods*, arXiv preprint arXiv:1902.08297 (2019).
- [13] Alfréd Rényi et al., *On measures of information and entropy*, Proceedings of the 4th berkeley symposium on mathematics, statistics and probability, 1961.

(Lexing Ying) DEPARTMENT OF MATHEMATICS, STANFORD UNIVERSITY, STANFORD, CA 94305
E-mail address: `lexing@stanford.edu`