

Approximately Gaussian Replicator Flows: Nonconvex Optimization as a Nash-Convergent Evolutionary Game

Brendon G. Anderson*

Samuel Pfrommer†

Somayeh Sojoudi†

Abstract—This work leverages tools from evolutionary game theory to solve unconstrained nonconvex optimization problems. Specifically, we lift such a problem to an optimization over probability measures, whose minimizers exactly correspond to the Nash equilibria of a particular population game. To algorithmically solve for such Nash equilibria, we introduce *approximately Gaussian replicator flows* (AGRFs) as a tractable alternative to simulating the corresponding infinite-dimensional replicator dynamics. Our proposed AGRF dynamics can be integrated using off-the-shelf ODE solvers when considering objectives with closed-form integrals against a Gaussian measure. We theoretically analyze AGRF dynamics by explicitly characterizing their trajectories and stability on quadratic objective functions, in addition to analyzing their descent properties. Our methods are supported by illustrative experiments on a range of canonical nonconvex optimization benchmark functions.

I. INTRODUCTION

We consider solving unconstrained optimization problems of the form

$$\inf_{x \in \mathbb{R}^n} f(x),$$

where f is generally nonconvex. Such optimization problems arise naturally across a range of scientific and engineering disciplines, including computational chemistry [8], computer vision [21], and medicine [13]. Notably, the classical empirical risk minimization approach in machine learning amounts to an unconstrained nonconvex optimization problem [19].

Evolutionary algorithms comprise one popular family of approaches for solving nonconvex problems. These algorithms stochastically search the parameter space by evolving a population of candidate solutions according to biologically-inspired variation and selection operations [29]. Unlike gradient-based methods like stochastic gradient descent, practical evolutionary algorithms are generally intractable for theoretical analysis due to the handcrafted nature of their update rules. Existing analyses are thus typically limited to a very simple (1+1)-evolution strategy [12, 27], or a discrete optimization setting [15]. Conversely, negative results show that evolutionary algorithms cannot converge exponentially fast for certain classes of objective functions [30].

Instead of relying on handcrafted evolution heuristics, we propose an evolutionarily-inspired optimization method that

naturally arises from a principled recasting of the optimization problem as a population game. In doing so, we achieve the following contributions:

- 1) We develop a framework for reformulating nonconvex optimization as a population game, which we show can be solved to global optimality by simulating the replicator dynamics for continuous objectives with unique global minima.
- 2) We propose *approximately Gaussian replicator flows* (AGRFs) as a tractable alternative to simulating the replicator dynamics, and theoretically analyze the resulting ordinary differential equation’s descent properties, in addition to determining its explicit solutions for quadratic objective functions.
- 3) We provide illustrative experiments using canonical nonconvex optimization test functions, demonstrating that our optimization method can escape local minima, adapt to the optimization landscape, and recover from poorly conditioned initializations.

A. Related Works

Our work, drawing on the frameworks and tools of population games and evolutionary dynamics [28], is intimately related to evolutionary algorithms for optimization. Some of the most classical of such methods are evolution strategies (ESs), wherein a population of candidate solutions iteratively evolves according to stochastic mutation and selection rules [26]. Covariance matrix adaptation evolution strategy (CMA-ES) extends ESs by allowing the covariance matrix of the ES sampling distribution to adapt to the geometry of the objective function [14]. CMA-ES is often considered the state-of-the-art in evolutionary computation [12], yet it has also been shown to be subsumed by natural evolution strategies (NESs) [1].

Unlike the heuristic approach to ES-based update rules, NES provides a principled optimization method by using natural gradient descent [2] to minimize the expected objective loss over a statistical manifold of parameterized probability distributions [31]. Closely related are estimation of distribution algorithms (EDAs) [20], which explicitly build, sample from, and update probability distributions to represent promising candidate solutions to the optimization, albeit these distributions are typically constructed via maximum likelihood estimation. Generalizing ES, CMA-ES, and NES, the information geometric optimization (IGO) framework poses optimization as a continuous time natural gradient flow on a statistical manifold [23]. IGO algorithms are constructed

*Department of Mechanical Engineering, California Polytechnic State University. Email: bga@calpoly.edu.

†Department of Electrical Engineering and Computer Sciences, University of California, Berkeley. Emails: {sam.pfrommer, sojoudi}@berkeley.edu.

This work was supported by grants from ONR, NSF, and ARO, and was conducted while Brendon G. Anderson was with the University of California, Berkeley.

by discretizing both time (via Euler's method) and spatial integration (via sampling) to approximate the flow.

All of the aforementioned evolutionary algorithms view the objective f as a black box, and as such, rely on evaluating f on a finite number of random samples. In practice, the challenges posed by the objective function f are often structural and not due to a lack of information, e.g., training machine learning models often relies on nonconvex and nonsmooth objectives [19, 10, 3], whose functional forms are known. Our framework is capable of leveraging closed-form integrals of nonconvex yet white-box objectives f against a Gaussian distribution, in order to pose evolutionary optimization as a *deterministic* differential equation. That is, unlike the above stochastic search methods, which often rely on heuristically designed algorithmic rules and parameters, our evolutionarily-inspired optimization framework can be methodically implemented via off-the-shelf ordinary differential equation (ODE) solvers without random sampling.

Two of the works closest related to ours are Jaćimović [17, 18]. The authors draw links between the replicator dynamics and minimization of a bilinear-quadratic objective function by showing that replicator trajectories are natural gradient flows for such objectives. This characterization relies on inverting the Fisher information matrix to evaluate the flow's vector field, whereas our approximately Gaussian replicator flows admit closed-form solutions for general quadratic functions. Our framework is particularly distinct due to the fact that our proposed AGRFs allow for using off-the-shelf ODE solvers whenever the objective is closed-form integrable against a Gaussian measure, e.g., general polynomials and sinusoids.

B. Notations

We denote the identity matrix by I , the dimension of which will always be clear from context. If X is a topological space, the set of finite signed Borel measures on X is denoted by $\mathcal{M}(X)$, and the set of Borel probability measures is denoted by $\mathcal{P}(X)$. The support of a measure $\mu \in \mathcal{M}(X)$ is denoted by $\text{supp}(\mu)$. For $\mu \in \mathcal{M}(X)$ and a measurable real-valued function f on X , we define $\langle f, \mu \rangle := \int_X f d\mu$, which we allow to be infinite. The derivative of a measure-valued mapping $\mu: [0, \infty) \rightarrow \mathcal{M}(X)$ at $t \in [0, \infty)$, if it exists, is the measure $\dot{\mu}(t) \in \mathcal{M}(X)$ defined by

$$\lim_{\epsilon \rightarrow 0} \left\| \frac{\mu(t+\epsilon) - \mu(t)}{\epsilon} - \dot{\mu}(t) \right\|_{\text{TV}} = 0,$$

where $\|\cdot\|_{\text{TV}}$ is the total variation norm on $\mathcal{M}(X)$. The Dirac distribution at $x \in X$ is denoted by δ_x , and the Gaussian distribution on \mathbb{R}^n with mean $m \in \mathbb{R}^n$ and covariance $C \in \mathbb{R}^{n \times n}$ is denoted by $\mathcal{N}(m, C)$.

II. OPTIMIZATION AS AN INFINITE-STRATEGY GAME

We seek to solve the unconstrained optimization problem

$$p^* := \inf_{x \in X} f(x), \quad (1)$$

over the Euclidean space $X = \mathbb{R}^n$ with a possibly nonsmooth and nonconvex objective function $f: X \rightarrow \mathbb{R}$. We assume that f is Borel measurable and attains a minimum.

The problem p^* is equivalent to an infinite-dimensional convex optimization problem over probability measures [4]:

$$p^* = p' := \inf_{\mu \in \mathcal{P}(X)} \langle f, \mu \rangle. \quad (2)$$

It is easy to see that p' also attains a minimizer, and that, if μ^* solves p' , then x^* solves p^* for all $x^* \in \text{supp}(\mu^*)$. Thus, our attention may be refocused on solving p' . We now show that solving p' amounts to finding a Nash equilibrium of a certain population game.

A. Turning Optimizers into Nash Equilibria

We briefly recall the key notions used in the study of population games and evolutionary dynamics—see Sandholm [28] for a more thorough treatment.

Consider a large population of strategic individuals playing a game. The players choose strategies from a *strategy set* S , which is assumed to be a topological space. A distribution of strategy choices employed across the population is encoded by a probability measure $\mu \in \mathcal{P}(S)$, termed a *population state*. Every population state has an associated *mean payoff function* $F_\mu: S \rightarrow \mathbb{R}$, assumed to be continuous, with $F_\mu(s)$ giving the payoff to strategy s when the population state is μ . The mapping $F: \mu \mapsto F_\mu$ is called the *population game*. A *Nash equilibrium* of the game F is a population state μ such that $\langle F_\mu, \nu \rangle \leq \langle F_\mu, \mu \rangle$ for all $\nu \in \mathcal{P}(S)$, and, if the inequality holds strictly for all $\nu \neq \mu$, then μ is called a *strict Nash equilibrium* of F . The set of all Nash equilibria of F is denoted by $\text{NE}(F)$.

Throughout the remainder of the paper, we consider a population game where the strategy set is X , and the mean payoff functions are given by *excess payoffs* under the objective function f :

$$F_\mu(x) = \langle f, \mu \rangle - f(x).$$

Such mean payoff functions can be written as the average of the *relative payoffs* $g(x, x') := f(x') - f(x)$ with respect to the population state μ :

$$F_\mu(x) = \int_X g(x, x') d\mu(x').$$

Note that this is an infinite-strategy population game since $X = \mathbb{R}^n$ is continuous. A key feature of this game is that it puts optimizers of p' in one-to-one correspondence with Nash equilibria:

Proposition 1. *A distribution $\mu \in \mathcal{P}(X)$ is a Nash equilibrium of F if and only if μ is a global minimizer of p' . Furthermore, μ is a strict Nash equilibrium of F if and only if μ is the unique global minimizer of p' .*

Proof. Let $\mu \in \text{NE}(F)$. Then, by definition of Nash equilibria, it holds for all $\nu \in \mathcal{P}(X)$ that

$$\begin{aligned} & \int_X \int_X (f(x') - f(x)) d\mu(x') d\nu(x) \\ & \leq \int_X \int_X (f(x') - f(x)) d\mu(x') d\mu(x), \end{aligned}$$

and hence

$$\langle f, \mu \rangle - \langle f, \nu \rangle \leq \langle f, \mu \rangle - \langle f, \mu \rangle = 0,$$

proving that $\langle f, \mu \rangle \leq \langle f, \nu \rangle$. Thus, μ solves p' .

On the other hand, suppose that $\mu \in \mathcal{P}(X)$ solves p' , so that $\langle f, \mu \rangle \leq \langle f, \nu \rangle$ for all $\nu \in \mathcal{P}(X)$. Then reversing the above line of analysis yields that $\langle F(\mu), \nu \rangle \leq \langle F(\mu), \mu \rangle$ for all $\nu \in \mathcal{P}(X)$, implying that $\mu \in \text{NE}(F)$. The second assertion follows by the same reasoning using the appropriate strict inequalities. \square

Another characteristic of this game is that it is *monotone*, meaning that $\langle F_\mu - F_\nu, \mu - \nu \rangle \leq 0$ for all $\mu, \nu \in \mathcal{P}(X)$:

Proposition 2. *The population game F is monotone.*

Proof. Let $\mu, \nu \in \mathcal{P}(X)$. Then

$$\begin{aligned} \langle F_\mu - F_\nu, \mu - \nu \rangle &= \int_X \int_X (f(x') - f(x)) d(\mu - \nu)(x') d(\mu - \nu)(x) \\ &= 0. \end{aligned}$$

\square

Monotonicity of population games has been found to be a useful condition for ensuring the convergence of evolutionary dynamics to Nash equilibria—see, e.g., [16, 11, 24, 6, 5].

B. Replicator Dynamics as an Optimizing Process

The *replicator dynamics* are one of the most popular evolutionary dynamics models (EDMs). EDMs are processes in which players continuously revise their strategies according to a certain revision protocol [28]. Specifically, the replicator dynamics model the situation in which players choose to imitate the strategy of a random opponent who currently has higher payoff, with probability proportional to the current difference in their payoffs. Mathematically, the replicator dynamics read

$$\dot{\mu}(t)(B) = \int_B (E(\delta_x, \mu(t)) - E(\mu(t), \mu(t))) d(\mu(t))(x) \quad (3)$$

for all Borel sets B and times $t \in [0, \infty)$, where $E: \mathcal{P}(X) \times \mathcal{P}(X) \rightarrow \mathbb{R}$, given by

$$E(\nu, \mu) = \int_X F_\mu d\nu = \langle f, \mu \rangle - \langle f, \nu \rangle,$$

encodes the *average mean payoff* to a population state $\nu \in \mathcal{P}(X)$ relative to a population state $\mu \in \mathcal{P}(X)$ according to the payoffs of our game F . Our key insight is the following: if the replicator dynamics converge to a Nash equilibrium of our game, then simulating them will lead us to a global minimizer of p' according to Proposition 1. Indeed, the following result shows that the monotonicity of our game ensures that this convergence happens in the case that f is continuous and has a unique global minimizer, even when f is nonconvex and nonsmooth.

Theorem 1. *Let $\mu_0 \in \mathcal{P}(X)$, and let $\mu: [0, \infty) \rightarrow \mathcal{P}(X)$ be a solution to the replicator dynamics with initial condition*

$\mu(0) = \mu_0$. If f is continuous and has a unique global minimizer $x^ \in X$, then it holds that $\mu(t)$ converges weakly to $\mu^* := \delta_{x^*}$ as $t \rightarrow \infty$.*

Proof. Suppose that f has a unique global minimizer $x^* \in X$. Clearly, $\mu^* = \delta_{x^*}$ is the unique global minimizer of p' . Thus, μ^* is a strict Nash equilibrium of F by Proposition 1 and hence $\langle F_{\mu^*}, \nu \rangle < \langle F_{\mu^*}, \mu^* \rangle$ for all $\nu \in \mathcal{P}(X) \setminus \{\mu^*\}$. By Proposition 2, it holds for all such ν that

$$\begin{aligned} \langle F_\nu, \nu \rangle - \langle F_\nu, \mu^* \rangle &= \langle F_{\mu^*}, \nu \rangle - \langle F_{\mu^*}, \mu^* \rangle + \langle F_\nu - F_{\mu^*}, \nu - \mu^* \rangle \\ &< 0. \end{aligned}$$

Thus, μ^* is strongly uninvadable with uniform invasion barrier $\epsilon = \infty$, per Oechssler and Riedel [22, Definition 5]. The result then follows from Oechssler and Riedel [22, Theorem 3]. \square

C. Simulating the Replicator Dynamics

Theorem 1 shows that we can find a globally optimal solution to (1) by simulating the replicator dynamics (3) with our constructed population game F . However, this constitutes solving an infinite-dimensional differential equation in general. It was shown in Cressman et al. [9] that, if $\mu: [0, \infty) \rightarrow \mathcal{P}(X)$ is a solution to the replicator dynamics, then its mean and covariance

$$\begin{aligned} m(t) &:= \int_X x d(\mu(t))(x), \\ C(t) &:= \int_X x x^\top d(\mu(t))(x) - m(t)m(t)^\top, \end{aligned}$$

evolve according to the finite-dimensional ordinary differential equations

$$\begin{aligned} \dot{m}_i(t) &= - \left. \frac{\partial}{\partial \lambda_i} E(\mu_\lambda(t), \mu(t)) \right|_{\lambda=0}, \\ \dot{C}_{ij}(t) &= \left. \frac{\partial^2}{\partial \lambda_i \partial \lambda_j} E(\mu_\lambda(t), \mu(t)) \right|_{\lambda=0}, \end{aligned}$$

for all $i, j \in \{1, \dots, n\}$, where $\mu_\lambda(t) \in \mathcal{P}(X)$ is the measure defined by

$$\mu_\lambda(t)(B) = \frac{\int_B \exp(-\lambda^\top x) d(\mu(t))(x)}{\int_X \exp(-\lambda^\top x) d(\mu(t))(x)}.$$

It is easy to show that these ODEs can be rewritten as

$$\begin{aligned} \dot{m}_i(t) &= m_i(t) \mathbb{E}_{x \sim \mu(t)}[f(x)] - \mathbb{E}_{x \sim \mu(t)}[x_i f(x)], \\ \dot{C}_{ij}(t) &= (C_{ij}(t) - m_i(t)m_j(t)) \mathbb{E}_{x \sim \mu(t)}[f(x)] \\ &\quad - \mathbb{E}_{x \sim \mu(t)}[x_i x_j f(x)] \\ &\quad + m_i(t) \mathbb{E}_{x \sim \mu(t)}[x_j f(x)] \\ &\quad + m_j(t) \mathbb{E}_{x \sim \mu(t)}[x_i f(x)]. \end{aligned} \quad (4)$$

Cressman et al. [9] showed that, in the case that the relative payoff function g takes a bilinear-quadratic form, the set of Gaussian distributions is invariant under (3), and therefore these two finite-dimensional ODEs fully characterize the replicator dynamics. In such cases, it holds that $C(t) \rightarrow 0$ as $t \rightarrow \infty$, and therefore the trajectory $t \mapsto \mu(t)$ converges

towards $t \mapsto \delta_{m(t)}$ (in the weak topology). Thus, simulating the above finite-dimensional ODEs to generate a limiting Dirac measure δ_{m^*} naturally gives us a minimizer m^* for this bilinear-quadratic case, through an application of Theorem 1.

In the general case, solving the ODEs (4) is intractable, as computing the expectations cannot be carried out for an arbitrary distribution $\mu(t)$ and objective function f . To make the differential equations solvable (using off-the-shelf ODE solvers), we propose to approximate $\mu(t)$ by $\mathcal{N}(m(t), C(t))$ at all times t . This leads to our primary definition:

Definition 1. A solution $(m, C): [0, \infty) \rightarrow \mathbb{R}^n \times \mathbb{R}^{n \times n}$ to the system of ordinary differential equations

$$\begin{aligned} \dot{m}_i(t) &= m_i(t) \mathbb{E}_{x \sim \mathcal{N}(m(t), C(t))} [f(x)] \\ &\quad - \mathbb{E}_{x \sim \mathcal{N}(m(t), C(t))} [x_i f(x)], \\ \dot{C}_{ij}(t) &= (C_{ij}(t) - m_i(t)m_j(t)) \mathbb{E}_{x \sim \mathcal{N}(m(t), C(t))} [f(x)] \\ &\quad - \mathbb{E}_{x \sim \mathcal{N}(m(t), C(t))} [x_i x_j f(x)] \\ &\quad + m_i(t) \mathbb{E}_{x \sim \mathcal{N}(m(t), C(t))} [x_j f(x)] \\ &\quad + m_j(t) \mathbb{E}_{x \sim \mathcal{N}(m(t), C(t))} [x_i f(x)], \end{aligned} \quad (5)$$

is called an *approximately Gaussian replicator flow (AGRF)* under f .

The right-hand sides of (5) can be determined explicitly for any objective functions f for which we can compute the appropriate Gaussian expectations in closed form. This class of objective functions, which we refer to as *Gaussian-integrable*, includes a wide variety of functions encountered in practice, as we will now discuss further.

D. Closed-Form AGRF Dynamics

Explicitly writing out the AGRF dynamics (5) requires computing the expectations of $f(x)$, $x_i f(x)$, and $x_i x_j f(x)$ under Gaussian-distributed vectors x . For general objectives f , these quantities can be approximated using Monte Carlo schemes. However, evaluating the desired estimates to sufficient resolution with such an approach is computationally intensive. We thus restrict our focus to Gaussian-integrable objective functions throughout the remainder of the paper, with polynomials and sinusoids serving as running examples.

1) *Polynomials*: For polynomials, the Gaussian expectations can be symbolically derived in closed-form from the following well-known identity [7]:

$$\mathbb{E}_{x \sim \mathcal{N}(0, I)} [x^\alpha] = \begin{cases} \pi^{-\frac{n}{2}} \prod_{i=1}^n 2^{\frac{\alpha_i}{2}} \Gamma\left(\frac{\alpha_i+1}{2}\right) & \alpha_i \text{ all even,} \\ 0 & \text{otherwise,} \end{cases}$$

where $x^\alpha = x_1^{\alpha_1} \dots x_n^{\alpha_n}$ is a monomial of order $\alpha = (\alpha_1, \dots, \alpha_n)$ and Γ is Euler's gamma function. This result is generalized to arbitrary mean and covariance by symbolically applying an affine transformation to x , and substituting for a particular mean m and covariance C . The extension from monomials to polynomials follows trivially from linearity of integration. Notably, since $x_i f(x)$ and $x_i x_j f(x)$ are all polynomials if $f(x)$ is a polynomial, we can integrate all these expressions using the same approach.

2) *Sinusoids*: For sinusoids (including cosine), we can derive closed-form expectations for complex exponentials $x \mapsto e^{ia^\top x}$ by computing the Fourier transform of multivariate Gaussian density functions. Straightforward calculus results in the following expressions, where, here, we use j, k as indices to reserve the symbol i for the imaginary unit:

$$\begin{aligned} \mathbb{E}_{x \sim \mathcal{N}(m, C)} [e^{ia^\top x}] &= e^{ia^\top m - \frac{1}{2} a^\top C a} =: \bar{m}, \\ \mathbb{E}_{x \sim \mathcal{N}(m, C)} [x_j e^{ia^\top x}] &= (m_j + i C_j^\top a) \bar{m}, \\ \mathbb{E}_{x \sim \mathcal{N}(m, C)} [x_j x_k e^{ia^\top x}] &= \left(C_{jk} + m_j m_k - (C_j^\top a)(C_k^\top a) \right. \\ &\quad \left. + i(m_j C_k^\top a + m_k C_j^\top a) \right) \bar{m}, \end{aligned}$$

where C_j is the j th row of C . The appropriate expectations for sinusoids are then derived using Euler's formula.

Remark 1. Our proposed approach for deterministically approximating the replicator dynamics using AGRFs intimately relies on the closed-form computation of the ODE derivatives, which requires Gaussian-integrable objective functions. We highlight two potential extensions to more complex objectives which are not as well-behaved. The first consists of estimating expectations using a Monte Carlo scheme and leveraging ODE integration methods which are better suited for noisy derivatives. The second avenue involves using existing approximation techniques to locally represent the objective in a Gaussian-integrable form, e.g., Taylor expansions for polynomials, or Fourier series for sinusoids. Given such a local approximation, the ODE derivatives are again evaluable in closed-form. Investigating these extensions is an exciting area of future work.

III. THEORETICAL ANALYSIS OF AGRFS

We now theoretically analyze our proposed approximately Gaussian replicator flows. Our analysis emphasizes AGRFs for quadratic objective functions, which serves to provide insight into the AGRF behavior in the vicinity of local minima of more general twice continuously differentiable functions. We begin by explicitly characterizing the governing equations of the AGRF for a quadratic objective function.

A. Explicit Characterization of Quadratic Flows

The following theorem is similar to Cressman et al. [9, Theorem 1], albeit their result is restricted to evolutionary games with bilinear-quadratic payoffs of the form $g(x, x') = -x^\top x + x^\top Q x'$, and thus cannot directly be used for our purposes where $g(x, x') = f(x') - f(x)$.

Theorem 2. Consider $f: \mathbb{R}^n \rightarrow \mathbb{R}$ given by $f(x) = x^\top A x + b^\top x + c$ with $A \in \mathbb{R}^{n \times n}$ symmetric. The AGRF under f is given by

$$\begin{aligned} \dot{m}(t) &= -2C(t)A m(t) - C(t)b \\ \dot{C}(t) &= -2C(t)A C(t). \end{aligned} \quad (6)$$

Proof. Let $i, j \in \{1, \dots, n\}$. Using the formulas for Gaussian expectations of linear, quadratic, and cubic forms from

Petersen et al. [25] and Jaćimović [18], we have that

$$\begin{aligned}
\dot{m}_i(t) &= m_i(t) \mathbb{E}_{x \sim \mathcal{N}(m(t), C(t))} [x^\top A x + b^\top x + c] \\
&\quad - \mathbb{E}_{x \sim \mathcal{N}(m(t), C(t))} [x_i (x^\top A x + b^\top x + c)] \\
&= m_i(t) (\text{tr}(AC(t)) + m(t)^\top A m(t) + b^\top m(t) + c) \\
&\quad - m_i(t) m(t)^\top A m(t) - 2e_i^\top C(t) A m(t) \\
&\quad - m_i(t) \text{tr}(AC(t)) - m_i(t) b^\top m(t) \\
&\quad - e_i^\top C(t) b - m_i(t) c \\
&= -2e_i^\top C(t) A m(t) - e_i^\top C(t) b.
\end{aligned}$$

We repeat a similar computation for the covariance matrix. For conciseness, we drop the time argument t for both C and m :

$$\begin{aligned}
\dot{C}_{ij} &= (C_{ij} - m_i m_j) \mathbb{E}_{x \sim \mathcal{N}(m, C)} [x^\top A x + b^\top x + c] \\
&\quad - \mathbb{E}_{x \sim \mathcal{N}(m, C)} [x_i x_j (x^\top A x + b^\top x + c)] \\
&\quad + m_i \mathbb{E}_{x \sim \mathcal{N}(m, C)} [x_j (x^\top A x + b^\top x + c)] \\
&\quad + m_j \mathbb{E}_{x \sim \mathcal{N}(m, C)} [x_i (x^\top A x + b^\top x + c)] \\
&= (C_{ij} - m_i m_j) (\text{tr}(AC) + m^\top A m + b^\top m + c) \\
&\quad - 2e_i^\top C A C e_j - 2m^\top A C (m_j e_i + m_i e_j) \\
&\quad - (\text{tr}(AC) + m^\top A m) (C_{ij} + m_i m_j) \\
&\quad - m_i m_j b^\top m - (m_i e_j^\top + m_j e_i^\top) C b - b^\top m C_{ij} \\
&\quad - c (m_i m_j + C_{ij}) \\
&\quad + m_i (m_j m^\top A m + 2e_j^\top C A m + m_j \text{tr}(AC) \\
&\quad + m_j b^\top m + e_j^\top C b + m_j c) \\
&\quad + m_j (m_i m^\top A m + 2e_i^\top C A m + m_i \text{tr}(AC) \\
&\quad + m_i b^\top m + e_i^\top C b + m_i c).
\end{aligned}$$

After simplification, we conclude that

$$\dot{C}_{ij}(t) = -2e_i^\top C(t) A C(t) e_j.$$

□

The ODE governing C in (6) is autonomous, and independent of m . Its solution is known in closed-form:

Proposition 3 (Jaćimović [17]). *Consider $f: \mathbb{R}^n \rightarrow \mathbb{R}$ given by $f(x) = x^\top A x + b^\top x + c$ with $A \in \mathbb{R}^{n \times n}$ symmetric. Then $C: [0, \infty) \rightarrow \mathbb{R}^{n \times n}$ given by*

$$C(t) = (C(0)^{-1} + 2tA)^{-1},$$

if it exists for all $t \in [0, \infty)$, solves the ODE $\dot{C}(t) = -2C(t)AC(t)$ of the AGRF under f given by (6).

The inverse matrix $(C(0)^{-1} + 2tA)^{-1}$ exists for all t whenever $C(0)$ is positive definite and A is positive semidefinite.

Next, we identify a closed-form solution to the remaining part of the AGRF for quadratic f , namely, the solution of the ODE governing m . Such a solution was not previously determined in Cressman et al. [9] or Jaćimović [17].

Proposition 4. *Consider $f: \mathbb{R}^n \rightarrow \mathbb{R}$ given by $f(x) = x^\top A x + b^\top x + c$ with $A \in \mathbb{R}^{n \times n}$ symmetric. Let $C: [0, \infty) \rightarrow$*

$\mathbb{R}^{n \times n}$ solve the ODE $\dot{C}(t) = -2C(t)AC(t)$ of the AGRF under f given by (6). Then $m: [0, \infty) \rightarrow \mathbb{R}^n$ given by

$$m(t) = C(t) (C(0)^{-1} m(0) - tb) \quad (7)$$

solves the ODE $\dot{m}(t) = -2C(t)Am(t) - C(t)b$ of the AGRF under f given by (6).

Proof. First, notice that the proposed solution m , given by (7), satisfies the required initial condition:

$$m(0) = C(0) (C(0)^{-1} m(0) - 0 \cdot b) = m(0).$$

Next, for the proposed solution m , we find that

$$\begin{aligned}
\dot{m}(t) &= \frac{d}{dt} (C(t) (C(0)^{-1} m(0) - tb)) \\
&= \dot{C}(t) (C(0)^{-1} m(0) - tb) - C(t)b \\
&= -2C(t)AC(t)(C(0)^{-1} m(0) - tb) - C(t)b \\
&= -2C(t)Am(t) - C(t)b,
\end{aligned}$$

where we used the fact that C solves $\dot{C}(t) = -2C(t)AC(t)$. Hence, m indeed solves the desired ODE. □

We now discuss the limiting behavior of AGRFs under quadratic objectives. Consider the mean trajectory (7). When A is positive definite, we have that

$$\begin{aligned}
\lim_{t \rightarrow \infty} (C(0)^{-1} + 2tA)^{-1} &= 0, \\
\lim_{t \rightarrow \infty} t(C(0)^{-1} + 2tA)^{-1} &= \frac{1}{2}A^{-1},
\end{aligned}$$

implying that m stabilizes to

$$\begin{aligned}
\lim_{t \rightarrow \infty} m(t) &= \lim_{t \rightarrow \infty} (C(0)^{-1} + 2tA)^{-1} (C(0)^{-1} m(0) - tb) \\
&= -\frac{1}{2}A^{-1}b.
\end{aligned}$$

This is precisely the solution to the optimization

$$\inf_{x \in \mathbb{R}^n} f(x) = \inf_{x \in \mathbb{R}^n} (x^\top A x + b^\top x + c),$$

so indeed, AGRFs exactly solve convex quadratic optimization problems.

Notice that one may write the dynamics for m as a preconditioned gradient flow:

$$\begin{aligned}
\dot{m}(t) &= - (C(0)^{-1} + 2tA)^{-1} (2Am(t) + b) \\
&= - (C(0)^{-1} + t\nabla^2 f(m(t)))^{-1} \nabla f(m(t)).
\end{aligned}$$

For small t , the dynamics resemble a gradient flow that is weighted according to the initial covariance at small times, representing the initial uncertainty in where a minimum resides. As time evolves, the flow begins to more closely resemble Newton's method. Jaćimović [17] showed that such replicator dynamics evolving under quadratic functions f and restricted to the statistical manifold of Gaussian distributions follow a natural gradient flow with respect to the Fisher information metric. However, they did not explicitly determine the mean trajectory m , nor was the flow's preconditioning matrix identified as having the form we determined above.

B. Descent of Flows

In this section, we study when AGRFs descend the objective's landscape. We begin by showing that descent always occurs for AGRFs on a convex quadratic objective.

Proposition 5. Consider $f: \mathbb{R}^n \rightarrow \mathbb{R}$ given by $f(x) = x^\top A x + b^\top x + c$ with $A \in \mathbb{R}^{n \times n}$ symmetric. If $C(0)$ is positive definite, A is positive semidefinite, and (m, C) is the AGRF under f , then the velocity $\dot{m}(t)$ is a descent direction of f for all $t \in [0, \infty)$, i.e.,

$$\frac{d}{dt} f(m(t)) = \nabla f(m(t))^\top \dot{m}(t) < 0.$$

Proof. Since $\nabla f(x) = 2Ax + b$, we find from Proposition 3 that

$$\begin{aligned} \nabla f(m(t))^\top \dot{m}(t) &= -(2Am(t) + b)(C(0)^{-1} + 2tA)^{-1}(2Am(t) + b) \\ &< 0 \end{aligned}$$

whenever $m(t)$ is not a minimizer of f . \square

Proposition 5 indicates that AGRF will converge to a local minimum of a general twice continuously differentiable function f when the bulk of the Gaussian measure $\mathcal{N}(m(t), C(t))$ is concentrated in a neighborhood around the minimum that locally looks like a strictly convex quadratic. However, such descent need not always occur, as we now illustrate.

Example 1. Consider $f: \mathbb{R} \rightarrow \mathbb{R}$ defined by $f(x) = \frac{3}{2}x^4 - \frac{1}{4}x^3 - 3x^2 + \frac{3}{4}x + 1$. It is easy to verify that f has one local minimum at $x = 1$ with $f(1) = 0$, a global minimum at $x = -1$ with $f(-1) = -1$, and a local maximum in between. A simple computation gives that the AGRF mean satisfies

$$\begin{aligned} \dot{m}(t) &= -C(t) \left(6m(t)^3 - \frac{3}{4}m(t)^2 - 6m(t) + \frac{3}{4} \right) \\ &\quad - C(t)^2 \left(18m(t) - \frac{3}{4} \right) \\ &= -C(t)f'(m(t)) - \frac{1}{2}C(t)^2 f'''(m(t)). \end{aligned}$$

The first term, $-C(t)f'(m(t))$, incentivizes the AGRF to perform local exploitation (descent along the gradient), whereas the second term, $-\frac{1}{2}C(t)^2 f'''(m(t))$, incentivizes the AGRF to explore for a better minimum. In particular, we may view the second term as descending along the gradient of f'' ; this term attempts to find a point where f has the steepest negative curvature, i.e., a sharp peak. With a large enough covariance $C(t)$, this term dominates and as such results in a hill-climbing effect, so that the flow can effectively look for better minima.

The above example shows that AGRFs are capable of escaping local minima in order to reach a global minimum. We empirically verify such beneficial behavior in Section IV.

We now discuss the descent behavior of AGRFs for more general functions f . For the remainder of this section, we will

assume that f has been adequately shifted so that $f(x) > 0$ for all $x \in X$, for convenience. The descent condition reads

$$\begin{aligned} \frac{d}{dt} f(m(t)) &= \nabla f(m(t))^\top \dot{m}(t) \\ &= \mathbb{E}_{x \sim \mathcal{N}(m(t), C(t))} [f(x) \nabla f(m(t))^\top (m(t) - x)] \\ &< 0. \end{aligned}$$

If the covariance $C(t)$ is small enough, then $x \sim \mathcal{N}(m(t), C(t))$ is close to $m(t)$ with high probability. For such points x that are close to $m(t)$, the first-order Taylor estimate $\nabla f(m(t))^\top (m(t) - x) \approx f(m(t)) - f(x)$ serves as a good approximation. Therefore, the above descent condition is well-approximated by

$$\mathbb{E}_{x \sim \mathcal{N}(m(t), C(t))} [f(x)(f(m(t)) - f(x))] < 0 \quad (8)$$

for small $C(t)$. The approximate descent condition (8) requires that, on average (weighted by the function f), it holds that $f(m(t)) < f(x)$, according to our current distribution of strategies $\mathcal{N}(m(t), C(t))$, which one may view as reflecting our “trust” in where the optimizer is. From this perspective, it makes sense for the dynamics to descend the objective, since it is believed that, on average, the current mean is a good one, and hence we should exploit it and trust a local descent. On the contrary, if $f(m(t)) > f(x)$ on average, then the approximate descent condition is violated. This indicates a situation in which our current distribution $\mathcal{N}(m(t), C(t))$ believes that the mean $m(t)$ performs worse than the average of the other strategies, and therefore ascent occurs in order to explore for alternatives.

We conclude this section by remarking that (nonconstant) convex functions f always satisfy the approximate descent condition (8), since

$$\begin{aligned} \mathbb{E}_{x \sim \mathcal{N}(m(t), C(t))} [f(x)f(m(t))] &- \mathbb{E}_{x \sim \mathcal{N}(m(t), C(t))} [f(x)^2] \\ &\leq \mathbb{E}_{x \sim \mathcal{N}(m(t), C(t))} [f(x)]^2 - \mathbb{E}_{x \sim \mathcal{N}(m(t), C(t))} [f(x)^2] \\ &< 0 \end{aligned}$$

by Jensen's inequality and the positivity of the variance of $f(x)$ under $x \sim \mathcal{N}(m(t), C(t))$.

IV. EXPERIMENTS

This section illustrates our methods with proof-of-concept nonconvex optimization experiments. We use canonical test functions, commonly used for evaluating optimization algorithms, to showcase the qualitative features of AGRFs.

Implementation. We compute mean and covariance time derivatives using the formulae in Section II-C. The resulting AGRF ODEs are integrated using an off-the-shelf Runge-Kutta 2-3 (RK23) solver. All optimizations are executed for a maximum of $T = 30$ seconds of simulated time, terminating early when the covariance matrix determinant is smaller than $\epsilon = 0.0001$. We emphasize that, unlike most other optimization algorithms, our framework requires no hyperparameter choices beyond the initial conditions $m(0), C(0)$ and the simulation stopping time T , which are standard inputs required by off-the-shelf ODE solvers.




Fig. 1: First two images visualize mean (black dots) and covariance (circles) versus ODE time (light gray is first solver iteration, black is last solver iteration). (a) Styblinski-Tang objective with a small initial covariance. (b) Styblinski-Tang with a large initial covariance. (c) Objective value with respect to ODE iteration for trajectory in (b).

A. Escaping Local Minima

Figure 1 illustrates the local minima-escaping behavior of our dynamics. We consider the Styblinski-Tang function, shifted to be more easily visualized on a log scale:

$$f(x) = 78.43 + \frac{1}{2} \left(\sum_{i=1}^n x_i^4 - 16x_i^2 + 5x_i \right).$$

We let $n = 2$, which gives rise to a global minimum at $(-2.903, -2.903)$ and three other local minima that are not global.

We fix an initial mean $m(0) = (3, 2)$ near the top-right local minimum. Figure 1a initializes $C(0) = 2I$, and Figure 1b initializes $C(0) = 30I$. Larger initial covariance intuitively helps the ODE “see” over the hump to eventually find the global minimum (Figure 1c). While the true replicator dynamics evolving over general probability measures would be guaranteed to find the global minimum by Theorem 1, we observe in Figure 1a that with a small initial covariance the ODE remains stuck in a local minimum. This arises from constraining our measures to be strictly Gaussian over the entire ODE trajectory.

Our framework naturally handles any Gaussian-integrable objective function, including the Styblinski-Tang polynomial objective. We show that sinusoids are also tractable by considering the classical Rastrigin function:

$$f(x) = 10n + \sum_{i=1}^n (x_i^2 - 10 \cos(2\pi x_i)).$$

As seen in Figure 2a, our algorithm replicates the local minima-escaping behavior observed above for a sufficiently large initial covariance of $C(0) = 10I$ and initial mean of $m(0) = (4, 4)$.

B. Optimization Landscape Adaptation

Figure 2b analyzes the tendency of the covariance dynamics to accelerate along directions of slow descent. Here, we consider the bivariate quadratic objective function given by

$$f(x) = (x_1 - 3)^2 + 4(x_2 - 3)^2,$$

which exhibits a smaller gradient along the x_1 -direction. The AGRF dynamics naturally stretch the covariance matrix

along this direction, accelerating convergence (Figure 2b). This is reminiscent of inverse Hessian preconditioning in Newton’s method, and the effect of momentum in stochastic gradient descent-based optimization.

C. Recovery from Poor Initialization

Figure 2c and Figure 2d illustrate the robustness of our approach to poorly conditioned initial points. The objective is a classic three-hump camel function:

$$f(x) = 2x_1^2 - 1.05x_1^4 + \frac{x_1^6}{6} + x_1x_2 + x_2^2.$$

The unique global minimum of this function lies at the origin, with a local minimum on each side. At points with sufficiently large x_1 -values, this function exhibits an ill-conditioned Hessian and a large-magnitude gradient that is nearly orthogonal to the x_2 -direction. We test such a poorly conditioned initial mean of $m(0) = (4, 4)$, with initial covariance matrices of $C(0) = 10I$ for Figure 2c and $C(0) = 100I$ for Figure 2d. In both cases, the ODE succeeds in recovering from the poor initialization and converges to the global minimum. This suggests that our approach successfully leverages the adaptive step size capabilities of the underlying RK23 ODE solver.

V. CONCLUSION

This work proposes a framework for recasting nonconvex optimization as an evolutionary game. Since simulating the resulting infinite-dimensional replicator dynamics is intractable, we propose to approximate the evolution by a trajectory of Gaussian distributions, termed an approximately Gaussian replicator flow. We theoretically characterize the flow’s descent behavior and ability to ascend from local minima, in addition to fully determining the flow on quadratic objectives. Our proof-of-concept experiments on polynomial and sinusoidal nonconvex benchmark objectives demonstrate the potential of our method for finding global optima using off-the-shelf ODE solvers, without the need for heuristic algorithm design choices. Our results demonstrate that even a Gaussian approximation to the replicator dynamics exhibits a range of desirable phenomenon, including the ability to escape local minima, adapt to the optimization landscape, and recover from poor initialization.




Fig. 2: Plots visualize mean (black dots) and covariance (circles) versus ODE time (light gray is first solver iteration, black is last solver iteration). (a) Rastrigin objective. (b) Quadratic objective $f(x) = (x_1 - 3)^2 + 4(x_2 - 3)^2$. (c) Three-hump objective with $m(0) = (4, 4)$ and $C(0) = 10I$. (d) Three-hump objective with $m(0) = (4, 4)$ and $C(0) = 100I$.

REFERENCES

- [1] Y. Akimoto, Y. Nagata, I. Ono, and S. Kobayashi. Theoretical foundation for CMA-ES from information geometry perspective. *Algorithmica*, 2012.
- [2] S.-I. Amari. Natural gradient works efficiently in learning. *Neural Computation*, 1998.
- [3] B. G. Anderson and S. Sojoudi. Global optimality guarantees for nonconvex unsupervised video segmentation. In *57th Annual Allerton Conference on Communication, Control, and Computing*, 2019.
- [4] B. G. Anderson, S. Pfrommer, and S. Sojoudi. Tight certified robustness via min-max representations of ReLU neural networks. In *IEEE Conference on Decision and Control*, 2023.
- [5] B. G. Anderson, S. Sojoudi, and M. Arcak. Evolutionary games on infinite strategy sets: Convergence to Nash equilibria via dissipativity. *arXiv preprint arXiv:2312.08286*, 2023.
- [6] M. Arcak and N. C. Martins. Dissipativity tools for convergence to Nash equilibria in population games. *IEEE Transactions on Control of Network Systems*, 2021.
- [7] A. Barvinok. Integration and optimization of multivariate polynomials by restriction onto a random subspace. *Foundations of Computational Mathematics*, 2007.
- [8] L. Chan, G. R. Hutchison, and G. M. Morris. Bayesian optimization for conformer generation. *Journal of Cheminformatics*, 2019.
- [9] R. Cressman, J. Hofbauer, and F. Riedel. Stability of the replicator equation for a single species with a multi-dimensional continuous trait space. *Journal of Theoretical Biology*, 2006.
- [10] S. Fattahi and S. Sojoudi. Exact guarantees on the absence of spurious local minima for non-negative rank-1 robust principal component analysis. *Journal of Machine Learning Research*, 2020.
- [11] M. J. Fox and J. S. Shamma. Population games, stable games, and passivity. *Games*, 2013.
- [12] T. Glasmachers. Global convergence of the $(1+1)$ evolution strategy to a critical point. *Evolutionary Computation*, 2020.
- [13] A. Gramfort, M. Kowalski, and M. Hämmäläinen. Mixed-norm estimates for the M/EEG inverse problem using accelerated gradient methods. *Physics in Medicine & Biology*, 2012.
- [14] N. Hansen and A. Ostermeier. Completely derandomized self-adaptation in evolution strategies. *Evolutionary Computation*, 2001.
- [15] J. He and G. Lin. Average convergence rate of evolutionary algorithms. *IEEE Transactions on Evolutionary Computation*, 2015.
- [16] J. Hofbauer and W. H. Sandholm. Stable games and their dynamics. *Journal of Economic Theory*, 2009.
- [17] V. Jaćimović. Natural gradient in evolutionary games. *arXiv preprint arXiv:2210.00573*, 2022.
- [18] V. Jaćimović. The fundamental theorem of natural selection in optimization and games. *Biosystems*, 2023.
- [19] P. Jain, P. Kar, et al. Non-convex optimization for machine learning. *Foundations and Trends in Machine Learning*, 2017.
- [20] P. Larrañaga and J. A. Lozano. *Estimation of distribution algorithms: A new tool for evolutionary computation*. Springer Science & Business Media, 2001.
- [21] Y. Ma, H. Mansour, D. Liu, P. T. Boufounos, and U. S. Kamilov. Accelerated image reconstruction for nonlinear diffractive imaging. In *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2018.
- [22] J. Oechssler and F. Riedel. Evolutionary dynamics on infinite strategy spaces. *Journal of Economic Theory*, 2001.
- [23] Y. Ollivier, L. Arnold, A. Auger, and N. Hansen. Information-geometric optimization algorithms: A unifying picture via invariance principles. *Journal of Machine Learning Research*, 2017.
- [24] S. Park, N. C. Martins, and J. S. Shamma. From population games to payoff dynamics models: A passivity-based approach. In *IEEE Conference on Decision and Control (CDC)*, 2019.
- [25] K. B. Petersen, M. S. Pedersen, et al. *The Matrix Cookbook*. Technical University of Denmark, 2008.
- [26] I. Rechenberg. Evolutionsstrategie. *Optimierung technischer Systeme nach Prinzipien der biologischen Evolution*, 1973.
- [27] G. Rudolph. Convergence rates of evolutionary algorithms for quadratic convex functions with rank-deficient Hessian. In *Adaptive and Natural Computing Algorithms: 11th International Conference*, 2013.
- [28] W. H. Sandholm. *Population Games and Evolutionary Dynamics*. MIT Press, 2010.
- [29] A. N. Sloss and S. Gustafson. 2019 evolutionary algorithms review. *Genetic Programming Theory and Practice*, 2020.
- [30] D. Tarłowski. On asymptotic convergence rate of evolutionary algorithms. *Authorea Preprints*, 2023.
- [31] D. Wierstra, T. Schaul, T. Glasmachers, Y. Sun, J. Peters, and J. Schmidhuber. Natural evolution strategies. *Journal of Machine Learning Research*, 2014.