Natural Gradient in Evolutionary Games

Vladimir Jaćimović

University of Montenegro, Faculty of Mathematics and Natural Sciences

FAU Erlangen, October 18th, 2022.

Continuous optimization: gradient-based methods

$$f(x) \to \max_{x \in \mathbb{R}^n}$$

Continuous optimization: gradient-based methods

$$f(x) \to \max_{x \in \mathbb{R}^n}$$

Gradient ascent:

$$x^{(t+\delta t)} = x^t + \delta t \nabla f(x^t). \quad (1)$$

Continuous optimization: gradient-based methods

$$f(x) \to \max_{x \in \mathbb{R}^n}$$

Gradient ascent:

$$x^{(t+\delta t)} = x^t + \delta t \nabla f(x^t). \tag{1}$$

In the limit of infinitesimally small step $\delta t \to 0$:

$$\frac{dx}{dt} = \nabla f(x) \Big|_{x=x(t)}. \tag{2}$$

Continuous optimization: gradient-based methods

$$f(x) \to \max_{x \in \mathbb{R}^n}$$

Gradient ascent:

$$x^{(t+\delta t)} = x^t + \delta t \nabla f(x^t). \tag{1}$$

In the limit of infinitesimally small step $\delta t \to 0$:

$$\frac{dx}{dt} = \nabla f(x)\Big|_{x=x(t)}. \tag{2}$$

Problems:

- Calculation of gradients can be difficult, or even unfeasible.
- The method will end up in a local maximum of f.

Another approach: Stochastic search

$$f(x) \to \max_{x \in \mathbb{R}^n}$$

Call f "the fitness function"

Another approach: Stochastic search

$$f(x) \to \max_{x \in \mathbb{R}^n}$$

Call f "the fitness function"

Another idea ("gradient-free optimization"):

Another approach: Stochastic search

$$f(x) \to \max_{x \in \mathbb{R}^n}$$

Call f "the fitness function"

Another idea ("gradient-free optimization"):
Pick n random points $x_1, \ldots, x_n$ in $\mathbb{R}^n$ and evaluate the fitness function in those points:

$$f(x_1), \ldots, f(x_n).$$

Then choose the best points (i.e. the maximal evaluation).

Another approach: Stochastic search

$$f(x) \to \max_{x \in \mathbb{R}^n}$$

Call f "the fitness function"

Another idea ("gradient-free optimization"):
Pick n random points $x_1, \dots, x_n$ in $\mathbb{R}^n$ and evaluate the fitness function in those points:

$$f(x_1), \dots, f(x_n).$$

Then choose the best points (i.e. the maximal evaluation).

Questions:

- What do you mean by "random"?

Another approach: Stochastic search

$$f(x) \to \max_{x \in \mathbb{R}^n}$$

Call f "the fitness function"

Another idea ("gradient-free optimization"):
Pick n random points $x_1, \ldots, x_n$ in $\mathbb{R}^n$ and evaluate the fitness function in those points:

$$f(x_1), \ldots, f(x_n).$$

Then choose the best points (i.e. the maximal evaluation).

Questions:

- What do you mean by "random"?
- Isn't is unwise to waste all the evaluations at once?

Another approach: Stochastic search

$$f(x) \to \max_{x \in \mathbb{R}^n}$$

Call f "the fitness function"

Another idea ("gradient-free optimization"):
Pick n random points $x_1, \ldots, x_n$ in $\mathbb{R}^n$ and evaluate the fitness function in those points:

$$f(x_1), \ldots, f(x_n).$$

Then choose the best points (i.e. the maximal evaluation).

Questions:

- What do you mean by "random"?
- Isn't is unwise to waste all the evaluations at once?

These two questions are the starting point for fields of stochastic search, evolutionary algorithms, genetic algorithms.

Stochastic search

Find a probability distribution $P \in \mathcal{P}$ with the p.d.f. $p(x)$, in such a way to maximize the expected fitness:

$$\mathbb{E}_{\mathcal{P}}[f(x)] = \int f(x)p(x)dx \to \max_{\mathcal{P}}. \tag{3}$$

Stochastic search

Find a probability distribution $P \in \mathcal{P}$ with the p.d.f. $p(x)$, in such a way to maximize the expected fitness:

$$\mathbb{E}_{\mathcal{P}}[f(x)] = \int f(x)p(x)dx \to \max_{\mathcal{P}}. \tag{3}$$

It may be difficult to optimize over the space of **all** probability distributions on $\mathbb{R}^n$. In practice, it makes sense to restrict to a certain family of probability distributions that depend on θ:

Stochastic search

Find a probability distribution $P \in \mathcal{P}$ with the p.d.f. $p(x)$, in such a way to maximize the expected fitness:

$$\mathbb{E}_{\mathcal{P}}[f(x)] = \int f(x)p(x)dx \to \max_{\mathcal{P}}. \tag{3}$$

It may be difficult to optimize over the space of **all** probability distributions on $\mathbb{R}^n$. In practice, it makes sense to restrict to a certain family of probability distributions that depend on θ:

$$J(\theta) = \mathbb{E}_{\theta}[f(x)] = \int f(x)p(x \mid \theta)dx \to \max_{\theta}. \tag{4}$$

Log-likelihood trick

Now, we want to minimize the function $J(\theta) = \mathbb{E}_\theta[f(x)]$ over the space of parameters.
Using the so-called *log-likelihood trick* one can write:

$$\nabla_\theta J(\theta) = \nabla_\theta \int f(x) p(x \mid \theta) dx = \int f(x) \nabla_\theta p(x \mid \theta) dx =$$

Log-likelihood trick

Now, we want to minimize the function $J(\theta) = \mathbb{E}_\theta[f(x)]$ over the space of parameters.

Using the so-called *log-likelihood trick* one can write:

$$\nabla_\theta J(\theta) = \nabla_\theta \int f(x)p(x \mid \theta)dx = \int f(x)\nabla_\theta p(x \mid \theta)dx =$$

$$\int f(x)\nabla_\theta p(x \mid \theta)\frac{p(x \mid \theta)}{p(x \mid \theta)}dx = \int [f(x)\nabla_\theta \log p(x \mid \theta)]p(x \mid \theta)dx =$$

Log-likelihood trick

Now, we want to minimize the function $J(\theta) = \mathbb{E}_\theta[f(x)]$ over the space of parameters.

Using the so-called *log-likelihood trick* one can write:

$$\nabla_\theta J(\theta) = \nabla_\theta \int f(x)p(x \,|\, \theta)dx = \int f(x)\nabla_\theta p(x \,|\, \theta)dx =$$

$$\int f(x)\nabla_\theta p(x \,|\, \theta)\frac{p(x \,|\, \theta)}{p(x \,|\, \theta)}dx = \int [f(x)\nabla_\theta \log p(x \,|\, \theta)]p(x \,|\, \theta)dx =$$

$$\mathbb{E}_\theta[f(x)\nabla_\theta \log p(x \,|\, \theta)].$$

Hence, the search gradient can be estimated from samples $x_1, \ldots, x_m$:

$$\nabla_\theta J(\theta) \approx \frac{1}{m}\sum_{i=1}^{m} f(x_i)\nabla_\theta \log p(x_i \,|\, \theta). \tag{5}$$

Gradient again!

We can apply gradient ascent to function $J(\theta)$:

$$\theta^{(t+\delta t)} = \theta^t + \delta t \nabla J(\theta^t). \tag{6}$$

Gradient again!

We can apply gradient ascent to function $J(\theta)$:

$$\theta^{(t+\delta t)} = \theta^t + \delta t \nabla J(\theta^t). \tag{6}$$

So, gradient ascent again!

Gradient again!

We can apply gradient ascent to function $J(\theta)$:

$$\theta^{(t+\delta t)} = \theta^t + \delta t \nabla J(\theta^t). \tag{6}$$

So, gradient ascent again! But, gradient of the different function now.

Gradient again!

We can apply gradient ascent to function $J(\theta)$:

$$\theta^{(t+\delta t)} = \theta^t + \delta t \nabla J(\theta^t). \tag{6}$$

So, gradient ascent again! But, gradient of the different function now.

In the limit of infinitesimally small step $\delta t \to 0$, this algorithm yields the gradient flow:

$$\frac{d\theta}{dt} = \nabla J(\theta)\Big|_{\theta=\theta(t)}. \tag{7}$$

Gradient depends on metric

We optimize over a family of probability measures, which are parametrized by θ.

Gradient depends on metric

We optimize over a family of probability measures, which are parametrized by θ.

But, gradient depends on the metric. In particular, $\nabla J(\theta)$ denotes the vector of partial derivatives:

$$\nabla J = \frac{\partial J}{\partial \theta} = \left(\frac{\partial J}{\partial \theta_1} \cdots\cdots \frac{\partial J}{\partial \theta_n} \right)^T.$$

This vector is gradient of J only if the metric is Euclidean with the standard inner product.

Gradient depends on metric

We optimize over a family of probability measures, which are parametrized by θ.

But, gradient depends on the metric. In particular, $\nabla J(\theta)$ denotes the vector of partial derivatives:

$$\nabla J = \frac{\partial J}{\partial \theta} = \left(\frac{\partial J}{\partial \theta_1} \ldots \ldots \frac{\partial J}{\partial \theta_n} \right)^T.$$

This vector is gradient of J only if the metric is Euclidean with the standard inner product.
If the metric is defined by a positive-definite matrix A, then the gradient with respect to that metric is given by:

$$A^{-1}\nabla$$

Gradient depends on metric

We optimize over a family of probability measures, which are parametrized by θ.

But, gradient depends on the metric. In particular, $\nabla J(\theta)$ denotes the vector of partial derivatives:

$$\nabla J = \frac{\partial J}{\partial \theta} = \left(\frac{\partial J}{\partial \theta_1} \dots \dots \frac{\partial J}{\partial \theta_n} \right)^T.$$

This vector is gradient of J only if the metric is Euclidean with the standard inner product.
If the metric is defined by a positive-definite matrix A, then the gradient with respect to that metric is given by:

$$A^{-1}\nabla$$

So, the question is: Do we equip spaces of probability distributions with the standard Euclidean metric?

An example of 1-dimensional Gaussian measures

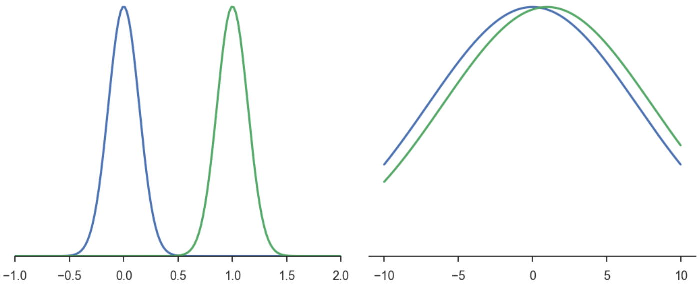

Left figure: Densities of $N(0, 0.2)$ and $N(1, 0.2)$. The Euclidean distance between them equals 1.

Right figure: Densities of $N(0, 10)$ and $N(1, 10)$. The Euclidean distance between them equals 1.

The Kullback-Leibler divergence

Euclidean distance on the space of Gaussians is inappropriate.

The Kullback-Leibler divergence

Euclidean distance on the space of Gaussians is inappropriate.

Let P and Q be two absolutely continuous distributions on $\mathbb{R}^n$ with p.d.f. $p(x)$ and $q(x)$, respectively.
The Kullback-Leibler divergence between P and Q is defined as:

$$I_{KL}(P \,||\, Q) = \int_{\mathbb{R}^n} p(x) \ln \left(\frac{p(x)}{q(x)} \right) dx.$$

K-L divergence is an appropriate measure how distant two random variables are.

The Kullback-Leibler divergence

Euclidean distance on the space of Gaussians is inappropriate.

Let P and Q be two absolutely continuous distributions on $\mathbb{R}^n$ with p.d.f. $p(x)$ and $q(x)$, respectively.
The Kullback-Leibler divergence between P and Q is defined as:

$$I_{KL}(P \,||\, Q) = \int_{\mathbb{R}^n} p(x) \ln\left(\frac{p(x)}{q(x)}\right) dx.$$

K-L divergence is an appropriate measure how distant two random variables are.
However, it is not the distance function.
Indeed, (like any divergence function) it is positive and satisfies $I_{KL}(P \,||\, Q) = 0 \iff P \equiv Q$.

The Kullback-Leibler divergence

Euclidean distance on the space of Gaussians is inappropriate.

Let P and Q be two absolutely continuous distributions on $\mathbb{R}^n$ with p.d.f. $p(x)$ and $q(x)$, respectively.
The Kullback-Leibler divergence between P and Q is defined as:

$$I_{KL}(P \,||\, Q) = \int_{\mathbb{R}^n} p(x) \ln \left(\frac{p(x)}{q(x)} \right) dx.$$

K-L divergence is an appropriate measure how distant two random variables are.
However, it is not the distance function.
Indeed, (like any divergence function) it is positive and satisfies $I_{KL}(P \,||\, Q) = 0 \iff P \equiv Q$.
But, K-L divergence is not symmetric, i.e. $I_{KL}(P \,||\, Q) \neq I_{KL}(Q \,||\, P)$.

The Kullback-Leibler divergence

But, locally, the K-L divergence is symmetric. More precisely, consider a family of probability distributions on $\mathbb{R}^n$. Denote this family by $\mathcal{P}$.
Suppose that this family depends on parameter θ.
Consider two distributions (random variables) $P(\theta + \delta\theta)$ and $P(\theta)$, where $\delta\theta$ is a small variation of the parameter.
Then: $I_{KL}(P(\theta + \delta\theta) \,||\, P(\theta)) \neq I_{KL}(P(\theta) \,||\, P(\theta + \delta\theta))$.

Statistical manifolds

Expanding the K-L divergence into Taylor series:

$$I_{KL}(P(\theta + \delta\theta) \,||\, P(\theta)) = 0 + 0 \cdot \delta\theta + \frac{1}{2}\delta\theta^T F(\theta)\delta\theta + o(\delta\theta^2), \qquad (8)$$

where $F(\theta)$ is the Hessian of the K-L divergence.

Statistical manifolds

Expanding the K-L divergence into Taylor series:

$$I_{KL}(P(\theta + \delta\theta) \,||\, P(\theta)) = 0 + 0 \cdot \delta\theta + \frac{1}{2}\delta\theta^T F(\theta)\delta\theta + o(\delta\theta^2), \qquad (8)$$

where $F(\theta)$ is the Hessian of the K-L divergence.

The matrix $F(\theta)$ is positive definite and hence defines a metric on the family $\mathcal{P}$.

Statistical manifolds

Expanding the K-L divergence into Taylor series:

$$I_{KL}(P(\theta + \delta\theta) \,||\, P(\theta)) = 0 + 0 \cdot \delta\theta + \frac{1}{2}\delta\theta^T F(\theta)\delta\theta + o(\delta\theta^2), \tag{8}$$

where $F(\theta)$ is the Hessian of the K-L divergence.

The matrix $F(\theta)$ is positive definite and hence defines a metric on the family $\mathcal{P}$.

Def: The metric defined by $F(\theta)$ is called the *Fisher information metric*.

Statistical manifolds

Expanding the K-L divergence into Taylor series:

$$I_{KL}(P(\theta+\delta\theta)\,||\,P(\theta)) = 0 + 0\cdot\delta\theta + \frac{1}{2}\delta\theta^T F(\theta)\delta\theta + o(\delta\theta^2), \tag{8}$$

where $F(\theta)$ is the Hessian of the K-L divergence.

The matrix $F(\theta)$ is positive definite and hence defines a metric on the family $\mathcal{P}$.

Def: The metric defined by $F(\theta)$ is called the *Fisher information metric*.

Hence, we have introduced a metric on the family of probability distributions, thus turning it into a Riemannian manifold.

Statistical manifolds

Expanding the K-L divergence into Taylor series:

$$I_{KL}(P(\theta+\delta\theta)\,||\,P(\theta)) = 0 + 0\cdot\delta\theta + \frac{1}{2}\delta\theta^T F(\theta)\delta\theta + o(\delta\theta^2), \qquad (8)$$

where $F(\theta)$ is the Hessian of the K-L divergence.

The matrix $F(\theta)$ is positive definite and hence defines a metric on the family $\mathcal{P}$.

Def: The metric defined by $F(\theta)$ is called the *Fisher information metric*.

Hence, we have introduced a metric on the family of probability distributions, thus turning it into a Riemannian manifold.

Such manifolds are called *statistical manifolds*.

Natural gradient

The above construction is the base of the field of Information Geometry. Information Geometry provides a universal way to equip families of probability distributions with Fisher information metric, thus turning them into Riemannian manifolds.

Natural gradient

The above construction is the base of the field of Information Geometry. Information Geometry provides a universal way to equip families of probability distributions with Fisher information metric, thus turning them into Riemannian manifolds.

Recall that we aim to maximize $J(\theta)$ over a statistical manifold $\mathcal{P}$, that is parameterized by θ.

Natural gradient

The above construction is the base of the field of Information Geometry. Information Geometry provides a universal way to equip families of probability distributions with Fisher information metric, thus turning them into Riemannian manifolds.

Recall that we aim to maximize $J(\theta)$ over a statistical manifold $\mathcal{P}$, that is parameterized by θ.

Taking into account the metric on $\mathcal{P}$, the gradient is given by:

$$\tilde{\nabla} = F^{-1}(\theta)\nabla.$$

Natural gradient

The above construction is the base of the field of Information Geometry. Information Geometry provides a universal way to equip families of probability distributions with Fisher information metric, thus turning them into Riemannian manifolds.

Recall that we aim to maximize $J(\theta)$ over a statistical manifold $\mathcal{P}$, that is parameterized by θ.

Taking into account the metric on $\mathcal{P}$, the gradient is given by:

$$\tilde{\nabla} = F^{-1}(\theta)\nabla.$$

$\tilde{\nabla}$ is called the *natural gradient*.

Natural gradient

The above construction is the base of the field of Information Geometry. Information Geometry provides a universal way to equip families of probability distributions with Fisher information metric, thus turning them into Riemannian manifolds.

Recall that we aim to maximize $J(\theta)$ over a statistical manifold $\mathcal{P}$, that is parameterized by θ.

Taking into account the metric on $\mathcal{P}$, the gradient is given by:

$$\tilde{\nabla} = F^{-1}(\theta)\nabla.$$

$\tilde{\nabla}$ is called the *natural gradient*.
So, natural gradient is a gradient on the statistical manifold w. r. to Fisher information metric.

Information-geometric optimization

Hence, the gradient ascent algorithm (6) is to be modified:

$$\theta^{(t+\delta t)} = \theta^t + \delta t \, F^{-1}(\theta^t) \nabla J(\theta^t).$$

Information-geometric optimization

Hence, the gradient ascent algorithm (6) is to be modified:

$$\theta^{(t+\delta t)} = \theta^t + \delta t \, F^{-1}(\theta^t) \nabla J(\theta^t).$$

Such an iterative procedure is called *natural gradient ascent*.

Information-geometric optimization

Hence, the gradient ascent algorithm (6) is to be modified:

$$\theta^{(t+\delta t)} = \theta^t + \delta t\, F^{-1}(\theta^t)\nabla J(\theta^t).$$

Such an iterative procedure is called *natural gradient ascent*.

The corresponding gradient flow is given by the following ODE

$$\frac{d\theta}{dt} = \tilde{\nabla} J(\theta)\Big|_{\theta=\theta(t)} = F^{-1}(\theta)\nabla J(\theta)\Big|_{\theta=\theta(t)}. \tag{9}$$

Now, we have an appropriate optimization procedure for the stochastic search.

Information-geometric optimization

Hence, the gradient ascent algorithm (6) is to be modified:

$$\theta^{(t+\delta t)} = \theta^t + \delta t \, F^{-1}(\theta^t) \nabla J(\theta^t).$$

Such an iterative procedure is called *natural gradient ascent*.

The corresponding gradient flow is given by the following ODE

$$\frac{d\theta}{dt} = \tilde{\nabla} J(\theta) \Big|_{\theta=\theta(t)} = F^{-1}(\theta) \nabla J(\theta) \Big|_{\theta=\theta(t)}. \tag{9}$$

Now, we have an appropriate optimization procedure for the stochastic search.
In general, we are free to choose any family of probability distributions on $\mathbb{R}^n$.

Information-geometric optimization

Hence, the gradient ascent algorithm (6) is to be modified:

$$\theta^{(t+\delta t)} = \theta^t + \delta t\, F^{-1}(\theta^t)\nabla J(\theta^t).$$

Such an iterative procedure is called *natural gradient ascent*.

The corresponding gradient flow is given by the following ODE

$$\frac{d\theta}{dt} = \tilde{\nabla} J(\theta)\Big|_{\theta=\theta(t)} = F^{-1}(\theta)\nabla J(\theta)\Big|_{\theta=\theta(t)}. \tag{9}$$

Now, we have an appropriate optimization procedure for the stochastic search.

In general, we are free to choose any family of probability distributions on $\mathbb{R}^n$. The main problem is that we have to evaluate (and invert) the Fisher information matrix for the chosen family.

The basic setup of Evolutionary Game Theory

Consider a population, consisting of n species:

$$A = \{A_1, \dots, A_n\}$$

The basic setup of Evolutionary Game Theory

Consider a population, consisting of n species:

$$A = \{A_1, \ldots, A_n\}$$

Denote by p_i a portion of species A_i in the population.
Denote by $\mathcal{P}(A)$ the set of all probability distributions over A. These distributions are usually called *categorical distributions*.

The basic setup of Evolutionary Game Theory

Consider a population, consisting of n species:

$$A = \{A_1, \ldots, A_n\}$$

Denote by p_i a portion of species A_i in the population.
Denote by $\mathcal{P}(A)$ the set of all probability distributions over A. These distributions are usually called *categorical distributions*.

The set $\mathcal{P}(A)$ can obviously be identified with the unit simplex:

$$\mathcal{P}(A) \simeq \Delta^n = \{p = (p_1 \ldots p_n)^T \ : \ p_1 + \cdots + p_n = 1, \ p_i \geq 0 \ \ i = 1, \ldots, n\}.$$

The basic setup of Evolutionary Game Theory

Consider a population, consisting of n species:

$$A = \{A_1, \dots, A_n\}$$

Denote by p_i a portion of species A_i in the population.
Denote by $\mathcal{P}(A)$ the set of all probability distributions over A. These distributions are usually called *categorical distributions*.

The set $\mathcal{P}(A)$ can obviously be identified with the unit simplex:

$$\mathcal{P}(A) \simeq \Delta^n = \{p = (p_1 \dots p_n)^T \ : \ p_1 + \cdots + p_n = 1, \ p_i \geq 0 \ \ i = 1, \dots, n\}.$$

To each species A_i we assign a certain continuous function $f_i(p_1, \dots, p_n)$.
We call f_i the *fitness function* of the species A_i.

The basic setup of Evolutionary Game Theory

Consider a population, consisting of n species:

$$A = \{A_1, \ldots, A_n\}$$

Denote by p_i a portion of species A_i in the population.
Denote by $\mathcal{P}(A)$ the set of all probability distributions over A. These distributions are usually called *categorical distributions*.

The set $\mathcal{P}(A)$ can obviously be identified with the unit simplex:

$$\mathcal{P}(A) \simeq \Delta^n = \{p = (p_1 \ldots p_n)^T \, : \, p_1 + \cdots + p_n = 1, \; p_i \geq 0 \;\; i = 1, \ldots, n\}.$$

To each species A_i we assign a certain continuous function $f_i(p_1, \ldots, p_n)$. We call f_i the *fitness function* of the species A_i.

Underline that fitness of species depends on portions of all species in the population.

Game-theoretic paradigm in study of the evolution

Hence, we have n fitness functions:
$f(p) = (f_1(p_1, \dots, p_n), \dots, f_n(p_1, \dots, p_n))^T$.
Vector-valued function $f(p)$ is called the *fitness landscape*.

Game-theoretic paradigm in study of the evolution

Hence, we have n fitness functions:
$f(p) = (f_1(p_1, \ldots, p_n), \ldots, f_n(p_1, \ldots, p_n))^T$.
Vector-valued function $f(p)$ is called the *fitness landscape*.

We refer to probability distributions from $\mathcal{P}(A)$ to as *strategies*.

Game-theoretic paradigm in study of the evolution

Hence, we have n fitness functions:
$f(p) = (f_1(p_1, \dots, p_n), \dots, f_n(p_1, \dots, p_n))^T$.
Vector-valued function $f(p)$ is called the *fitness landscape*.

We refer to probability distributions from $\mathcal{P}(A)$ to as *strategies*.

The strategy is said to be *pure*, if $p_i = 1$ and $p_j = 0$, $j \neq i$.

Game-theoretic paradigm in study of the evolution

Hence, we have n fitness functions:
$f(p) = (f_1(p_1, \dots, p_n), \dots, f_n(p_1, \dots, p_n))^T$.
Vector-valued function $f(p)$ is called the *fitness landscape*.

We refer to probability distributions from $\mathcal{P}(A)$ to as *strategies*.

The strategy is said to be *pure*, if $p_i = 1$ and $p_j = 0$, $j \neq i$.
All other strategies are said to be *mixed strategies*.

Some illustrative examples of evolutionary games

- The hawk-dove game
- The prisoner's dilemma
- Rock-scissors-paper
- Sharing of a cake
- Battle of sexes

Some illustrative examples of evolutionary games

- The hawk-dove game
- The prisoner's dilemma
- Rock-scissors-paper
- Sharing of a cake
- Battle of sexes

The evolution is not exclusively biological concept, one can talk about:

- cultural evolution
- behavioral evolution
- evolution of moral norms
- ...many more $\Longrightarrow$ a variety of interpretations of evolutionary games

Replicator equations

Introduce the mean fitness in the population:

$$\langle f(p) \rangle = p \cdot f(p) = \sum_{i=1}^{n} p_i f_i(p). \tag{10}$$

The notion $x \cdot y$ stands for the inner product of vectors x and y in the vector space.

Replicator equations

Introduce the mean fitness in the population:

$$\langle f(p) \rangle = p \cdot f(p) = \sum_{i=1}^{n} p_i f_i(p). \tag{10}$$

The notion $x \cdot y$ stands for the inner product of vectors x and y in the vector space.

Assume that portions p_i evolve in time, i.e. $p_i(t)$, $i = 1, \dots, n$. Moreover, $p_i(t)$ satisfy the following equations:

$$\dot{p}_i = p_i \left(f_i(p) - \langle f(p) \rangle \right), \quad i = 1, \dots, n. \tag{11}$$

Replicator equations

Introduce the mean fitness in the population:

$$\langle f(p)\rangle = p \cdot f(p) = \sum_{i=1}^{n} p_i f_i(p). \tag{10}$$

The notion $x \cdot y$ stands for the inner product of vectors x and y in the vector space.

Assume that portions p_i evolve in time, i.e. $p_i(t)$, $i = 1, \ldots, n$.
Moreover, $p_i(t)$ satisfy the following equations:

$$\dot{p}_i = p_i \left(f_i(p) - \langle f(p)\rangle \right), \quad i = 1, \ldots, n. \tag{11}$$

Equations (11) are called *replicator equations*.

Replicator equations

Introduce the mean fitness in the population:

$$\langle f(p) \rangle = p \cdot f(p) = \sum_{i=1}^{n} p_i f_i(p). \tag{10}$$

The notion $x \cdot y$ stands for the inner product of vectors x and y in the vector space.

Assume that portions p_i evolve in time, i.e. $p_i(t)$, $i = 1, \dots, n$.
Moreover, $p_i(t)$ satisfy the following equations:

$$\dot{p}_i = p_i \left(f_i(p) - \langle f(p) \rangle \right), \quad i = 1, \dots, n. \tag{11}$$

Equations (11) are called *replicator equations*.
They are intuitively clear.
We say that the state of population at moment t is given by vector $p(t) = (p_1(t), \dots, p_n(t))$. Equations (11) claim that if for a given state of population, species A_i has higher fitness than the mean fitness, its portion $p_i(t)$ will increase.

The K-L divergence on the set of categorical distributions

Notice that the unit simplex is an invariant set for the dynamics (11).

The K-L divergence on the set of categorical distributions

Notice that the unit simplex is an invariant set for the dynamics (11).

Hence, replicator equations describe the evolution of categorical probability distributions.

The K-L divergence on the set of categorical distributions

Notice that the unit simplex is an invariant set for the dynamics (11).

Hence, replicator equations describe the evolution of categorical probability distributions.

Question: How should we introduce metric on the set $\mathcal{P}(A)$?

The K-L divergence on the set of categorical distributions

Notice that the unit simplex is an invariant set for the dynamics (11).

Hence, replicator equations describe the evolution of categorical probability distributions.

Question: How should we introduce metric on the set $\mathcal{P}(A)$?

Given two probability distributions $p, q \in \mathcal{P}(A)$, the Kullback-Leibler divergence (sometimes also referred to as *relative entropy*) between them is defined as:

$$I_{KL}(p \,||\, q) = \sum_{i=1}^{n} q_i \ln \left(\frac{q_i}{p_i} \right). \tag{12}$$

The Shahshahani metric on the set of categorical distributions

The Taylor expansion of the K-L divergence up to the second order term along the diagonal $p = q$ yields:

$$I_{KL}(p \,||\, q) = I_{KL}(p \,||\, q)|_{p=q} + (\nabla I_{KL}(p \,||\, q)|_{p=q}) \cdot (p - q) +$$

$$+ \frac{1}{2}(p - q)^T H(p)(p - q) + \cdots = \tag{13}$$

$$= 0 + 0 + \frac{1}{2}(p - q)^T H(p)(p - q) + \cdots .$$

The Shahshahani metric on the set of categorical distributions

The Taylor expansion of the K-L divergence up to the second order term along the diagonal $p = q$ yields:

$$I_{KL}(p\,||\,q) = I_{KL}(p\,||\,q)|_{p=q} + (\nabla I_{KL}(p\,||\,q)|_{p=q}) \cdot (p-q) +$$
$$+\tfrac{1}{2}(p-q)^T H(p)(p-q) + \cdots = \tag{13}$$
$$= 0 + 0 + \tfrac{1}{2}(p-q)^T H(p)(p-q) + \cdots .$$

Here, $H(p)$ denotes the Hessian matrix of I_{KL}:

$$H(p) = \left(\frac{\partial^2 I_{KL}}{\partial p_i \partial q_j}\right)\Bigg|_{p=q}.$$

An easy calculation yields:

$$H(p) = diag\{1/p_1, \ldots, 1/p_n\}.$$

The Shahshahani metric

Hence, the K-L divergence on $\mathcal{P}(A)$ generates the metric given by the following positive-definite matrix

$$H(p) = diag\{1/p_1, \dots, 1/p_n\}$$

The Shahshahani metric

Hence, the K-L divergence on $\mathcal{P}(A)$ generates the metric given by the following positive-definite matrix

$$H(p) = diag\{1/p_1, \ldots, 1/p_n\}$$

This metric is called the *Shahshahani metric* in EGT.

The Shahshahani metric

Hence, the K-L divergence on $\mathcal{P}(A)$ generates the metric given by the following positive-definite matrix

$$H(p) = diag\{1/p_1, \dots, 1/p_n\}$$

This metric is called the *Shahshahani metric* in EGT.
We see that the Shahshahani metric is in fact the Fisher information metric for the particular case of the family $\mathcal{P}(A)$.

The Shahshahani metric

Hence, the K-L divergence on $\mathcal{P}(A)$ generates the metric given by the following positive-definite matrix

$$H(p) = diag\{1/p_1, \dots, 1/p_n\}$$

This metric is called the *Shahshahani metric* in EGT.
We see that the Shahshahani metric is in fact the Fisher information metric for the particular case of the family $\mathcal{P}(A)$.

In whole, we equipped $\mathcal{P}(A)$ with the Fisher information metric, thus turning it into a **statistical manifold**.

Gradient flows on the manifold $\mathcal{P}(A)$

Theorem 1. (Marc Harper) If the system of differential equations $\dot{p}_i = f_i(p)$ defines a Euclidean gradient flow with $f_i(p) = \frac{\partial V}{\partial p_i}$, then replicator equations (11) define a gradient flow with respect to the Shahshahani metric.

Gradient flows on the manifold $\mathcal{P}(A)$

Theorem 1. (Marc Harper) If the system of differential equations $\dot{p}_i = f_i(p)$ defines a Euclidean gradient flow with $f_i(p) = \frac{\partial V}{\partial p_i}$, then replicator equations (11) define a gradient flow with respect to the Shahshahani metric.

Assume that fitness functions are linear:

$$f_i(p) = \sum_{j=1}^{n} a_{ij} p_j, \text{ for some coefficients } a_{ij}.$$

Gradient flows on the manifold $\mathcal{P}(A)$

Theorem 1. (Marc Harper) If the system of differential equations $\dot{p}_i = f_i(p)$ defines a Euclidean gradient flow with $f_i(p) = \frac{\partial V}{\partial p_i}$, then replicator equations (11) define a gradient flow with respect to the Shahshahani metric.

Assume that fitness functions are linear:

$$f_i(p) = \sum_{j=1}^{n} a_{ij} p_j, \text{ for some coefficients } a_{ij}.$$

This assumption yields the linear fitness landscape $f(p) = Ap$ for matrix $A = \{a_{ij}\}$ which is called *fitness (or payoff) matrix*.

Gradient flows on the manifold $\mathcal{P}(A)$

Theorem 1. (Marc Harper) If the system of differential equations $\dot{p}_i = f_i(p)$ defines a Euclidean gradient flow with $f_i(p) = \frac{\partial V}{\partial p_i}$, then replicator equations (11) define a gradient flow with respect to the Shahshahani metric.

Assume that fitness functions are linear:

$$f_i(p) = \sum_{j=1}^{n} a_{ij} p_j, \text{ for some coefficients } a_{ij}.$$

This assumption yields the linear fitness landscape $f(p) = Ap$ for matrix $A = \{a_{ij}\}$ which is called *fitness (or payoff) matrix*.

For linear fitness functions Theorem 1 becomes even more transparent.

Gradient flows on the manifold $\mathcal{P}(A)$

Theorem 1. (Marc Harper) If the system of differential equations $\dot{p}_i = f_i(p)$ defines a Euclidean gradient flow with $f_i(p) = \frac{\partial V}{\partial p_i}$, then replicator equations (11) define a gradient flow with respect to the Shahshahani metric.

Assume that fitness functions are linear:

$$f_i(p) = \sum_{j=1}^{n} a_{ij} p_j, \text{ for some coefficients } a_{ij}.$$

This assumption yields the linear fitness landscape $f(p) = Ap$ for matrix $A = \{a_{ij}\}$ which is called *fitness (or payoff) matrix*.

For linear fitness functions Theorem 1 becomes even more transparent. Indeed, Euclidean gradient is Ap and Shahshahani potential turns out to be the mean fitness: $\frac{1}{2} p \cdot f(p) = \frac{1}{2} p \cdot Ap$.

Evolutionary dynamics as a Bayesian learning

Following the pioneering works of Maynard Smith, one of central concepts in EGT is the notion of an *evolutionary stable strategy* (ESS).

Evolutionary dynamics as a Bayesian learning

Following the pioneering works of Maynard Smith, one of central concepts in EGT is the notion of an *evolutionary stable strategy* (ESS).
A strategy $\hat{p} \in \mathcal{P}(A)$ is said to be ESS if $\hat{p} \cdot f(p) > p \cdot f(p)$ for all p in some neighborhood of $\hat{p}$.

Evolutionary dynamics as a Bayesian learning

Following the pioneering works of Maynard Smith, one of central concepts in EGT is the notion of an *evolutionary stable strategy* (ESS).
A strategy $\hat{p} \in \mathcal{P}(A)$ is said to be ESS if $\hat{p} \cdot f(p) > p \cdot f(p)$ for all p in some neighborhood of $\hat{p}$.

Theorem 2.(Marc Harper) The strategy $\hat{p}$ is an interior ESS, if and only if $I_{KL}(\hat{p} \,||\, p)$ is a Lyapunov function for the dynamics (11).

Evolutionary dynamics as a Bayesian learning

Following the pioneering works of Maynard Smith, one of central concepts in EGT is the notion of an *evolutionary stable strategy* (ESS).
A strategy $\hat{p} \in \mathcal{P}(A)$ is said to be ESS if $\hat{p} \cdot f(p) > p \cdot f(p)$ for all p in some neighborhood of $\hat{p}$.

Theorem 2.(Marc Harper) The strategy $\hat{p}$ is an interior ESS, if and only if $I_{KL}(\hat{p} \,||\, p)$ is a Lyapunov function for the dynamics (11).

Theorems 1 and 2 unveil an information-theoretic background of evolutionary games with a finite strategy set.

Evolutionary dynamics as a Bayesian learning

Following the pioneering works of Maynard Smith, one of central concepts in EGT is the notion of an *evolutionary stable strategy* (ESS).
A strategy $\hat{p} \in \mathcal{P}(A)$ is said to be ESS if $\hat{p} \cdot f(p) > p \cdot f(p)$ for all p in some neighborhood of $\hat{p}$.

Theorem 2.(Marc Harper) The strategy $\hat{p}$ is an interior ESS, if and only if $I_{KL}(\hat{p} \,||\, p)$ is a Lyapunov function for the dynamics (11).

Theorems 1 and 2 unveil an information-theoretic background of evolutionary games with a finite strategy set. Evolutionary dynamics appears as a learning process.

Evolutionary dynamics as a Bayesian learning

Following the pioneering works of Maynard Smith, one of central concepts in EGT is the notion of an *evolutionary stable strategy* (ESS).
A strategy $\hat{p} \in \mathcal{P}(A)$ is said to be ESS if $\hat{p} \cdot f(p) > p \cdot f(p)$ for all p in some neighborhood of $\hat{p}$.

Theorem 2.(Marc Harper) The strategy $\hat{p}$ is an interior ESS, if and only if $I_{KL}(\hat{p} \,||\, p)$ is a Lyapunov function for the dynamics (11).

Theorems 1 and 2 unveil an information-theoretic background of evolutionary games with a finite strategy set. Evolutionary dynamics appears as a learning process. K-L divergence can be seen as the amount of information left to learn until the population achieves ESS.

Evolutionary dynamics as a Bayesian learning

Following the pioneering works of Maynard Smith, one of central concepts in EGT is the notion of an *evolutionary stable strategy* (ESS).
A strategy $\hat{p} \in \mathcal{P}(A)$ is said to be ESS if $\hat{p} \cdot f(p) > p \cdot f(p)$ for all p in some neighborhood of $\hat{p}$.

Theorem 2.(Marc Harper) The strategy $\hat{p}$ is an interior ESS, if and only if $I_{KL}(\hat{p} \,||\, p)$ is a Lyapunov function for the dynamics (11).

Theorems 1 and 2 unveil an information-theoretic background of evolutionary games with a finite strategy set. Evolutionary dynamics appears as a learning process. K-L divergence can be seen as the amount of information left to learn until the population achieves ESS. For a nice and enlighting exposition of the above interpretations see the paper "The Relative Entropy in Biological Systems" by Baez and Pollard.

Information-geometric approach to evolutionary games

We have seen that the evolutionary dynamics can be seen as a flow on a statistical manifold.
This point of view introduces the setup of Information Geometry to Evolutionary Game Theory.

Information-geometric approach to evolutionary games

We have seen that the evolutionary dynamics can be seen as a flow on a statistical manifold.
This point of view introduces the setup of Information Geometry to Evolutionary Game Theory.

So, why not to study games in which strategies are members of a certain family of absolutely continuous probability measures?

Information-geometric approach to evolutionary games

We have seen that the evolutionary dynamics can be seen as a flow on a statistical manifold.
This point of view introduces the setup of Information Geometry to Evolutionary Game Theory.

So, why not to study games in which strategies are members of a certain family of absolutely continuous probability measures?
For instance, take the family of Gaussian measures $\mathcal{N}(a, C)$,

Information-geometric approach to evolutionary games

We have seen that the evolutionary dynamics can be seen as a flow on a statistical manifold.
This point of view introduces the setup of Information Geometry to Evolutionary Game Theory.

So, why not to study games in which strategies are members of a certain family of absolutely continuous probability measures?
For instance, take the family of Gaussian measures $\mathcal{N}(a, C)$, endow it with the Fisher information metric

Information-geometric approach to evolutionary games

We have seen that the evolutionary dynamics can be seen as a flow on a statistical manifold.
This point of view introduces the setup of Information Geometry to Evolutionary Game Theory.

So, why not to study games in which strategies are members of a certain family of absolutely continuous probability measures?
For instance, take the family of Gaussian measures $\mathcal{N}(a, C)$, endow it with the Fisher information metric thus turning it into the statistical manifold

Information-geometric approach to evolutionary games

We have seen that the evolutionary dynamics can be seen as a flow on a statistical manifold.
This point of view introduces the setup of Information Geometry to Evolutionary Game Theory.

So, why not to study games in which strategies are members of a certain family of absolutely continuous probability measures?
For instance, take the family of Gaussian measures $\mathcal{N}(a, C)$, endow it with the Fisher information metric thus turning it into the statistical manifold where the manifold coordinates are the mean vector a and covariance matrix C.

Information-geometric approach to evolutionary games

We have seen that the evolutionary dynamics can be seen as a flow on a statistical manifold.
This point of view introduces the setup of Information Geometry to Evolutionary Game Theory.

So, why not to study games in which strategies are members of a certain family of absolutely continuous probability measures?
For instance, take the family of Gaussian measures $\mathcal{N}(a, C)$, endow it with the Fisher information metric thus turning it into the statistical manifold where the manifold coordinates are the mean vector a and covariance matrix C.

Or, some other family of probability measures, for instance, Cauchy distributions, or...

Information-geometric approach to evolutionary games

We have seen that the evolutionary dynamics can be seen as a flow on a statistical manifold.
This point of view introduces the setup of Information Geometry to Evolutionary Game Theory.

So, why not to study games in which strategies are members of a certain family of absolutely continuous probability measures?
For instance, take the family of Gaussian measures $\mathcal{N}(a, C)$, endow it with the Fisher information metric thus turning it into the statistical manifold where the manifold coordinates are the mean vector a and covariance matrix C.

Or, some other family of probability measures, for instance, Cauchy distributions, or...

In such a setup the biological interpretation is slightly different. The points in a continuous space correspond to *traits*.

Evolutionary game with a continuous trait space

Consider a game in which Strategy of the population at the moment t is given by a probability measure $P(t)$ on a measurable space $(S, \mathcal{B})$, where $\mathcal{B}$ is the σ-algebra of Borel subsets of S.

Evolutionary game with a continuous trait space

Consider a game in which Strategy of the population at the moment t is given by a probability measure $P(t)$ on a measurable space $(S, \mathcal{B})$, where $\mathcal{B}$ is the σ-algebra of Borel subsets of S.

We will also say that the *population is in the state* $P(t)$ at the moment t. Denote by $\Delta(S)$ the set of all probability measures with respect to $(S, \mathcal{B})$.

Evolutionary game with a continuous trait space

Consider a game in which Strategy of the population at the moment t is given by a probability measure $P(t)$ on a measurable space $(S, \mathcal{B})$, where $\mathcal{B}$ is the σ-algebra of Borel subsets of S.

We will also say that the *population is in the state* $P(t)$ at the moment t. Denote by $\Delta(S)$ the set of all probability measures with respect to $(S, \mathcal{B})$.

For $A \in \mathcal{B}$,

$$P(t)(A) = \int_A P(t) ds$$

is the proportion of traits (or actions) belonging to the set A in the set of all traits (actions) at the moment t.

Evolutionary game with a continuous trait space

Consider a game in which Strategy of the population at the moment t is given by a probability measure $P(t)$ on a measurable space $(S, \mathcal{B})$, where $\mathcal{B}$ is the σ-algebra of Borel subsets of S.

We will also say that the *population is in the state* $P(t)$ at the moment t. Denote by $\Delta(S)$ the set of all probability measures with respect to $(S, \mathcal{B})$.

For $A \in \mathcal{B}$,

$$P(t)(A) = \int_A P(t) ds$$

is the proportion of traits (or actions) belonging to the set A in the set of all traits (actions) at the moment t.

The fitness landscape is given by a continuous real-valued function $f : S \times S \to \mathbb{R}^n$.

Expected fitness in evolutionary game with a continuous trait space

Then the fitness (expected payoff) of the strategy $Q \in \Delta(S)$ played against population in the state P is

$$\pi(Q, P) = \int_S \int_S f(s, y) Q(ds) P(dy). \tag{14}$$

Expected fitness in evolutionary game with a continuous trait space

Then the fitness (expected payoff) of the strategy $Q \in \Delta(S)$ played against population in the state P is

$$\pi(Q,P) = \int_S \int_S f(s,y) Q(ds) P(dy). \tag{14}$$

A trait $s \in S$ is identified with the delta distribution $\delta_s \in \Delta(S)$ that assigns probability one to s. These distributions are pure strategies.

Expected fitness in evolutionary game with a continuous trait space

Then the fitness (expected payoff) of the strategy $Q \in \Delta(S)$ played against population in the state P is

$$\pi(Q, P) = \int_S \int_S f(s, y) Q(ds) P(dy). \tag{14}$$

A trait $s \in S$ is identified with the delta distribution $\delta_s \in \Delta(S)$ that assigns probability one to s. These distributions are pure strategies. From (14) we have that the fitness of a trait s when the population is in the state P is given by:

$$\pi(s, P) \equiv \pi(\delta_s, P) = \int_S \int_S f(s, y) \delta_s(ds) P(dy) = \int_S f(s, y) P(dy).$$

Expected fitness in evolutionary game with a continuous trait space

Then the fitness (expected payoff) of the strategy $Q \in \Delta(S)$ played against population in the state P is

$$\pi(Q, P) = \int_S \int_S f(s, y) Q(ds) P(dy). \tag{14}$$

A trait $s \in S$ is identified with the delta distribution $\delta_s \in \Delta(S)$ that assigns probability one to s. These distributions are pure strategies. From (14) we have that the fitness of a trait s when the population is in the state P is given by:

$$\pi(s, P) \equiv \pi(\delta_s, P) = \int_S \int_S f(s, y) \delta_s(ds) P(dy) = \int_S f(s, y) P(dy).$$

Integration over S yields an expected fitness in the population

$$\langle \pi(P) \rangle \equiv \pi(P, P) = \int_S f(\delta_s, P) P(ds). \tag{15}$$

Replicator equations in evolutionary games with a continuous trait space

Finally, relative fitness of a trait s against the population in the state P is

$$\phi(s, P) = \pi(s, P) - \langle \pi(P) \rangle.$$

Replicator equations in evolutionary games with a continuous trait space

Finally, relative fitness of a trait s against the population in the state P is

$$\phi(s, P) = \pi(s, P) - \langle \pi(P) \rangle.$$

The replicator equation for the above game reads:

$$\frac{dP}{dt}(A) = \int_A \phi(s, P) P(ds) = \int_A (\pi(\delta_s) - \pi(P, P)) P(ds). \tag{16}$$

Replicator equations in evolutionary games with a continuous trait space

Finally, relative fitness of a trait s against the population in the state P is

$$\phi(s, P) = \pi(s, P) - \langle \pi(P) \rangle.$$

The replicator equation for the above game reads:

$$\frac{dP}{dt}(A) = \int_A \phi(s, P) P(ds) = \int_A (\pi(\delta_s) - \pi(P, P)) P(ds). \qquad (16)$$

Equation (16) has the same meaning as replicator equations (11): the probability of traits belonging to the set A increases if they have higher expected fitness than the mean fitness in the population. Hence, under dynamics (16) the probability measure $P(t)$ at each moment t tends to concentrate around those subsets of S in which traits have high fitness.

Quadratic-bilinear fitness landscape

Replicator equation (16) yields an infinite-dimensional dynamical system.
Very difficult to analyze, or simulate.

Quadratic-bilinear fitness landscape

Replicator equation (16) yields an infinite-dimensional dynamical system. Very difficult to analyze, or simulate.

In order to simplify the problem introduce two assumptions:

Quadratic-bilinear fitness landscape

Replicator equation (16) yields an infinite-dimensional dynamical system. Very difficult to analyze, or simulate.

In order to simplify the problem introduce two assumptions:

Assumption 1: The fitness landscape f is given by a bilinear-quadratic function:

$$f(s, y) = -s \cdot Qs + s \cdot By, \tag{17}$$

where Q is a positive symmetric $n \times n$ matrix and B is an arbitrary $n \times n$ matrix.

Quadratic-bilinear fitness landscape

Replicator equation (16) yields an infinite-dimensional dynamical system. Very difficult to analyze, or simulate.

In order to simplify the problem introduce two assumptions:

Assumption 1: The fitness landscape f is given by a bilinear-quadratic function:

$$f(s, y) = -s \cdot Qs + s \cdot By, \tag{17}$$

where Q is a positive symmetric $n \times n$ matrix and B is an arbitrary $n \times n$ matrix.

Assumption 2: The initial state of population $P(0) = P_0$ is given by a multivariate Gaussian distribution $N(a, C)$, with mean vector $m \in \mathbb{R}^n$ and positive-definite covariance matrix $C \in \mathbb{R}^n \times \mathbb{R}^n$.

Replicator equations restricted to the family of Gaussian measures

Theorem. (Cressman, Hofbauer, Riedel)
Under assumptions 1 and 2 the Gaussian family of measures is forward-invariant for the dynamics (16).
Moreover, assume that the initial distribution by $P(0) = N(a(0), C(0))$ is Gaussian with the mean vector $m(0)$ and covariance matrix $C(0)$. Then, the solution of (16) is given by $P(t) = N(a(t), C(t))$, with parameters $a(t)$ and $C(t)$ satisfying the following system

$$\frac{da(t)}{dt} = C(t)(B - 2Q)a(t);$$

$$\frac{dC(t)}{dt} = -2C(t)QC(t).$$

Maximization of mean fitness

Consider the problem of maximization of the mean fitness as defined in (15):

$$\langle \pi(P) \rangle = \mathbb{E}_P[f(s, y)] \to \max_{P \in \mathcal{N}(a,C)} \tag{18}$$

In other words, we are looking for the Gaussian probability measure that maximizes the mean fitness (expected payoff) in the population.

Maximization of mean fitness

Consider the problem of maximization of the mean fitness as defined in (15):

$$\langle \pi(P) \rangle = \mathbb{E}_P[f(s, y)] \to \max_{P \in \mathcal{N}(a, C)} \tag{18}$$

In other words, we are looking for the Gaussian probability measure that maximizes the mean fitness (expected payoff) in the population.

Taking into account Assumption 1, (18) can be rewritten as:

$$J(a, C) = \mathbb{E}_{a,C}[f(s, y)] = \mathbb{E}_{a,C}[-s \cdot Qs + s \cdot By] \to \max_{a,C} \tag{19}$$

Maximization of mean fitness

Consider the problem of maximization of the mean fitness as defined in (15):

$$\langle \pi(P) \rangle = \mathbb{E}_P[f(s, y)] \to \max_{P \in \mathcal{N}(a,C)} \tag{18}$$

In other words, we are looking for the Gaussian probability measure that maximizes the mean fitness (expected payoff) in the population.

Taking into account Assumption 1, (18) can be rewritten as:

$$J(a, C) = \mathbb{E}_{a,C}[f(s, y)] = \mathbb{E}_{a,C}[-s \cdot Qs + s \cdot By] \to \max_{a,C} \tag{19}$$

Evaluating mathematical expectation of the quadratic form and taking into account that s and y are mutually independent random variables, distributed as $N(a, C)$, we get

$$J(a, C) = \mathbb{E}_{a,C}[s \cdot Qs + s \cdot By] = -a \cdot Qa - Tr(QC) + a \cdot Ba \to \max_{a,C} .$$

Gradient of the mean fitness

Gradient of the above function is:

$$\nabla_{(a,C)} J(a,C) = \begin{pmatrix} \nabla_a J(a,C) \\ \nabla_C J(a,C) \end{pmatrix} = \begin{pmatrix} -2Qa + Ba \\ -Q \end{pmatrix}. \tag{20}$$

Gradient of the mean fitness

Gradient of the above function is:

$$\nabla_{(a,C)}J(a,C) = \begin{pmatrix} \nabla_a J(a,C) \\ \nabla_C J(a,C) \end{pmatrix} = \begin{pmatrix} -2Qa + Ba \\ -Q \end{pmatrix}. \tag{20}$$

Underline that this is "vanilla" gradient.
In order to get the natural gradient, we must evaluate the Fisher information matrix $F^{-1}(a,C)$.

Gradient of the mean fitness

Gradient of the above function is:

$$\nabla_{(a,C)} J(a,C) = \begin{pmatrix} \nabla_a J(a,C) \\ \nabla_C J(a,C) \end{pmatrix} = \begin{pmatrix} -2Qa + Ba \\ -Q \end{pmatrix}. \qquad (20)$$

Underline that this is "vanilla" gradient.
In order to get the natural gradient, we must evaluate the Fisher information matrix $F^{-1}(a,C)$.

Fortunately, this can be done, since we are restricted to the family of Gaussian measures. Then, natural gradient is $\tilde{\nabla}_{(a,C)} = F^{-1}(a,C)\nabla_{(a,C)}$.

Gradient of the mean fitness

Gradient of the above function is:

$$\nabla_{(a,C)} J(a,C) = \left(\begin{array}{c} \nabla_a J(a,C) \\ \nabla_C J(a,C) \end{array} \right) = \left(\begin{array}{c} -2Qa + Ba \\ -Q \end{array} \right). \tag{20}$$

Underline that this is "vanilla" gradient.
In order to get the natural gradient, we must evaluate the Fisher information matrix $F^{-1}(a,C)$.

Fortunately, this can be done, since we are restricted to the family of Gaussian measures. Then, natural gradient is $\tilde{\nabla}_{(a,C)} = F^{-1}(a,C)\nabla_{(a,C)}$.

Multiplying the vanilla gradient with $F^{-1}(a,C)$, we get

$$\tilde{\nabla}_{(a,C)} J(a,C) = \left(\begin{array}{c} C\nabla_a J(a,C) \\ 2C\nabla_C J(a,C)C \end{array} \right) = \left(\begin{array}{c} C(B-2Q)a \\ -2CQC \end{array} \right).$$

Natural gradient flow for maximization of the mean fitness

Hence, the natural gradient flow for maximization of the mean fitness is

$$\frac{da(t)}{dt} = C(t)(B - 2Q)a(t);$$
$$\frac{dC(t)}{dt} = -2C(t)QC(t).$$

Natural gradient flow for maximization of the mean fitness

Hence, the natural gradient flow for maximization of the mean fitness is

$$\frac{da(t)}{dt} = C(t)(B - 2Q)a(t);$$
$$\frac{dC(t)}{dt} = -2C(t)QC(t).$$

This is precisely the same system as replicator equations restricted to the manifold of Gaussian measures.
Hence, we have obtained the following

Natural gradient flow for maximization of the mean fitness

Hence, the natural gradient flow for maximization of the mean fitness is

$$\frac{da(t)}{dt} = C(t)(B - 2Q)a(t);$$
$$\frac{dC(t)}{dt} = -2C(t)QC(t).$$

This is precisely the same system as replicator equations restricted to the manifold of Gaussian measures.
Hence, we have obtained the following

Theorem.
The replicator equations (16) under assumptions 1 and 2 are the natural gradient flow for maximization of the mean fitness in population.

Natural Evolution Strategies

Hence, we have found that the replicator dynamics, restricted to the manifold of Gaussians is the natural gradient flow.

Natural Evolution Strategies

Hence, we have found that the replicator dynamics, restricted to the manifold of Gaussians is the natural gradient flow.

In order to explore some properties of natural gradient flows, consider the following maximization problem:

$$\begin{aligned} J(\theta + \delta\theta) - J(\theta) \to \max_{\delta\theta} \\ \text{so that } I_{KL}(p(\theta + \delta\theta) \,||\, p(\theta)) = \varepsilon. \end{aligned} \tag{21}$$

Natural Evolution Strategies

Hence, we have found that the replicator dynamics, restricted to the manifold of Gaussians is the natural gradient flow.

In order to explore some properties of natural gradient flows, consider the following maximization problem:

$$\begin{aligned} J(\theta+\delta\theta) - J(\theta) &\to \max_{\delta\theta} \\ \text{so that } I_{KL}(p(\theta+\delta\theta)\,||\,p(\theta)) &= \varepsilon. \end{aligned} \tag{21}$$

Hence, we are looking for an update $\delta\theta$ in the direction that yields a maximal increase of the expected fitness, while imposing constraint on the information gain at each step.

Natural Evolution Strategies

$$J(\theta + \delta\theta) - J(\theta) \to \max_{\delta\theta}$$
$$\text{so that } I_{KL}(p(\theta + \delta\theta) \,||\, p(\theta)) = \varepsilon. \tag{22}$$

Natural Evolution Strategies

$$
\begin{aligned}
J(\theta + \delta\theta) - J(\theta) \to \max_{\delta\theta} \\
\text{so that } I_{KL}(p(\theta + \delta\theta) \,||\, p(\theta)) = \varepsilon.
\end{aligned}
\tag{22}
$$

In order to solve (21), expand the objective function:

$$J(\theta + \delta\theta) - J(\theta) = \nabla J(\theta) \cdot \delta\theta + \cdots$$

and introduce the Lagrange function (using (8) and neglecting higher order terms)

$$L(\delta\theta, \lambda) = \nabla J(\theta) \cdot \delta\theta + \lambda(\varepsilon - \frac{1}{2}\delta\theta \cdot F(\theta)\delta\theta).$$

Natural Evolution Strategies

$$J(\theta + \delta\theta) - J(\theta) \to \max_{\delta\theta}$$
$$\text{so that } I_{KL}(p(\theta + \delta\theta) \,||\, p(\theta)) = \varepsilon. \tag{22}$$

In order to solve (21), expand the objective function:

$$J(\theta + \delta\theta) - J(\theta) = \nabla J(\theta) \cdot \delta\theta + \cdots$$

and introduce the Lagrange function (using (8) and neglecting higher order terms)

$$L(\delta\theta, \lambda) = \nabla J(\theta) \cdot \delta\theta + \lambda(\varepsilon - \frac{1}{2}\delta\theta \cdot F(\theta)\delta\theta).$$

Taking derivatives with respect to $\delta\theta$ and λ:

$$\frac{\partial L}{\partial \delta\theta} = \nabla J(\theta) - \lambda F(\theta)\,\delta\theta; \quad \frac{\partial L}{\partial \lambda} = \varepsilon - \frac{1}{2}\delta\theta \cdot F(\theta)\delta\theta.$$

Natural Evolution Strategies

Equating the derivative w.r. to θ to zero, one gets $\nabla J(\theta) - \lambda F \delta\theta = 0$. Solving for $\delta\theta$ yields

$$\delta\theta = \frac{1}{\lambda} F^{-1}(\theta) \nabla J(\theta).$$

By substituting $\delta\theta$ and expanding the K-L divergence, one obtains expressions for ε and λ. In the limit of infinitesimally small time increment this yields the following ODE:

$$\frac{d\theta}{dt} = F^{-1}(\theta) \nabla J(\theta) \Big|_{\theta=\theta(t)} \tag{23}$$

which is precisely the natural gradient flow for maximization of $J(\theta)$.

Evolutionary dynamics as Natural Evolution Strategies

Theorem.
The natural gradient ascent algorithm follows the direction $\delta\theta$ on statistical manifold $\mathcal{P}$ that achieves a maximal increase in expected fitness $J(\theta)$, for a given K-L divergence between random variables $P(\theta + \delta\theta)$ and $P(\theta)$.

Evolutionary dynamics as Natural Evolution Strategies

Theorem.
The natural gradient ascent algorithm follows the direction $\delta\theta$ on statistical manifold $\mathcal{P}$ that achieves a maximal increase in expected fitness $J(\theta)$, for a given K-L divergence between random variables $P(\theta + \delta\theta)$ and $P(\theta)$.

The black-box optimization community has adopted the term *Natural Evolution Strategies* for these algorithms based on the natural gradient.

Evolutionary dynamics as Natural Evolution Strategies

Theorem.
The natural gradient ascent algorithm follows the direction $\delta\theta$ on statistical manifold $\mathcal{P}$ that achieves a maximal increase in expected fitness $J(\theta)$, for a given K-L divergence between random variables $P(\theta + \delta\theta)$ and $P(\theta)$.

The black-box optimization community has adopted the term *Natural Evolution Strategies* for these algorithms based on the natural gradient.

Interpretation: Underline that $I_{KL}(p(\theta + \delta\theta) \,||\, p(\theta))$ can be interpreted as a loss of diversity when the distribution is varied by $\delta\theta$.

Evolutionary dynamics as Natural Evolution Strategies

Theorem.
The natural gradient ascent algorithm follows the direction $\delta\theta$ on statistical manifold $\mathcal{P}$ that achieves a maximal increase in expected fitness $J(\theta)$, for a given K-L divergence between random variables $P(\theta + \delta\theta)$ and $P(\theta)$.

The black-box optimization community has adopted the term *Natural Evolution Strategies* for these algorithms based on the natural gradient.

Interpretation: Underline that $I_{KL}(p(\theta + \delta\theta) \,||\, p(\theta))$ can be interpreted as a loss of diversity when the distribution is varied by $\delta\theta$.

Conclusion: The population in an evolutionary game follows the Natural Evolution Strategy. It tends to maximize its mean fitness, while constraining loss of the diversity.

Conclusions

Replicator equations are natural gradient flows for maximization of the mean fitness on statistical manifolds.

Conclusions

Replicator equations are natural gradient flows for maximization of the mean fitness on statistical manifolds.

The above assertion has been rigorously proven for evolutionary dynamics on two statistical manifolds:
a) manifold of categorical distributions with parameters $p_1, \dots, p_n$ (evolutionary games with a finite strategy set);

Conclusions

Replicator equations are natural gradient flows for maximization of the mean fitness on statistical manifolds.

The above assertion has been rigorously proven for evolutionary dynamics on two statistical manifolds:
a) manifold of categorical distributions with parameters $p_1, \dots, p_n$ (evolutionary games with a finite strategy set);
b) manifold of multivariate Gaussian distributions $\mathcal{N}(a, C)$.

Conclusions

Replicator equations are natural gradient flows for maximization of the mean fitness on statistical manifolds.

The above assertion has been rigorously proven for evolutionary dynamics on two statistical manifolds:
a) manifold of categorical distributions with parameters $p_1, \ldots, p_n$ (evolutionary games with a finite strategy set);
b) manifold of multivariate Gaussian distributions $\mathcal{N}(a, C)$.

It remains an open (and challenging) question if similar assertion is valid for evolutionary dynamics on other statistical manifolds.

Conclusion: asymptotic behavior of evolutionary dynamics

Contrary to what has frequently been claimed in the literature, algorithms based on natural gradient do not exhibit fast convergence.

Conclusion: asymptotic behavior of evolutionary dynamics

Contrary to what has frequently been claimed in the literature, algorithms based on natural gradient do not exhibit fast convergence.
This is not surprising, if we know that they correspond to Natural Evolution Strategies. Hence, they realize a compromise between maximization of the mean fitness and preservation of the diversity.

In many cases this can be an advantage (from the point of view of exploration-exploitation algorithms).

Applications

This offers new insights in those fields where evolutionary games serve as models:
"Macrofoundations" of EGT: collective ratio of the population.

Applications

This offers new insights in those fields where evolutionary games serve as models:
"Macrofoundations" of EGT: collective ratio of the population.

We have pointed out that evolutionary games yield natural gradient policies.
On the other side, natural gradient has originally been proposed by Amari with a view on algorithms of AI and ML.
Since then, the link between natural gradient optimization and Bayesian inference has been recognized. Natural gradient policies have been applied to training different models of NN's, such as VAE, GAN's, etc.

Applications

This offers new insights in those fields where evolutionary games serve as models:
"Macrofoundations" of EGT: collective ratio of the population.

We have pointed out that evolutionary games yield natural gradient policies.
On the other side, natural gradient has originally been proposed by Amari with a view on algorithms of AI and ML.
Since then, the link between natural gradient optimization and Bayesian inference has been recognized. Natural gradient policies have been applied to training different models of NN's, such as VAE, GAN's, etc.
We have demonstrated that these algorithms can be implemented through evolutionary games.

Application: Evolutionary games for multi-agent reinforcement learning

It has been recognized that multi-agent reinforcement learning algorithms are in essence evolutionary games (for games with a finite strategy set).

Application: Evolutionary games for multi-agent reinforcement learning

It has been recognized that multi-agent reinforcement learning algorithms are in essence evolutionary games (for games with a finite strategy set).

Our findings open a perspective for implementing MARL policies through simulating evolutionary games with a continuous trait space (at least, on the manifold of Gaussian distributions).

About EGT in general

J. Maynard Smith: "Evolution and the Theory of Games", Cambridge University Press, 1982.

J. McKenzie Alexander, Evolutionary Game Theory, in: The Stanford Encyclopedia of Philosophy (Summer 2021 Edition), Edward N. Zalta (ed.), URL = https://plato.stanford.edu/archives/sum2021/entries/game-evolutionary/

J. Hofbauer, P. Schuster and K. Sigmund, A note on evolutionary stable strategies and game dynamics, J. Theor. Biol., **81**(3) (1979), 609-612.

J. Hofbauer and K. Sigmund: "Evolutionary Games and Population Dynamics", Cambridge University Press, 1998.

W. H. Sandholm: "Population Games and Evolutionary Dynamics", MIT Press, 2010.

About evolutionary games with a continuous trait space

R. Cressman, J. Hofbauer and F. Riedel, Stability of the replicator equation for a single species with a multi-dimensional continuous trait space, J. Theor. Biol., **239**(2) (2006).

G. P. Karev and A. S. Novozhilov, How trait distributions evolve in populations with parametric heterogeneity, Math. Biosci., **315** (2019).

J. Oechssler and F. Riedel, Evolutionary Dynamics on Infinite Strategy Spaces, Econ. Theory, **17** (2001), 141-162.

M. Ruijgrok and T. W. Ruijgrok, An effective replicator equation for games with a continuous strategy set, Dyn. Games Appl., **5** (2015), 157-179.

About Information Geometry and natural gradient

S. Amari, Natural gradient works efficiently in learning, Neur. Comp. **10** (2) (1998), 251-276.

S. Amari and H. Nagaoka, "Methods of Information Geometry" American Mathematical Society, Providence, RL, 2007.

A. Fujiwara and S. Amari, Gradient system in view of information geometry, Physica D: Nonlinear Phenomena, **80** (3) (1995), 317-327.

About geometric approach to EGT

E. Akin: "The Geometry of Population Genetics", Springer, Berlin, 1979.

J. C. Baez and B. S. Pollard, Relative entropy in biological systems, Entropy, **18** (2), 46 (2016).

M. Harper, Information geometry and evolutionary game theory, (2009) arXiv:0911.1383.

M. Harper, The replicator equation as inference dynamics, (2009) arXiv:0911.1763.

S. Shahshahani: "A new mathematical framework for the study of linkage and selection", Memoirs of the AMS, **17**, 221, 1979.

About black-box optimization and Natural Evolution Strategies

Y. Akimoto, Y. Nagata, I. Ono and S. Kobayashi, Theoretical Foundation for CMA-ES from Information Geometry Perspective, Algorithmica, **64** (2012), 698-716.

H.-G. Beyer, Convergence analysis of evolutionary algorithms that are based on paradigms of information geometry, Evol. Comput., **22**(4) (2014), 679-709.

T. Glasmachers, Convergence of the IGO-Flow of Isotropic Gaussian Distributions on Convex Quadratic Problems, in: Parallel Problem Solving from Nature - PPSN XII. PPSN 2012. Lecture Notes in Computer Science, vol 7491. Springer, Berlin, Heidelberg (2012)

About black-box optimization and Natural Evolution Strategies

N. Hansen and A. Auger, Principled Design of Continuous Stochastic Search: From Theory to Practice, in: Theory and Principled Methods for Designing Metaheustics, Y. Borenstein and A. Moraglio (eds) Springer, (2014) 145-180.

N. Hansen and A. Ostermeier, Completely derandomized self-adaptation in evolution strategies, Evol. Comput., **9**(2)(2001), 159-195.

L. Malagó and G. Pistone, Information Geometry of the Gaussian Distribution in View of Stochastic Optimization, in: FOGA '15: Proceedings of the 2015 ACM Conference on Foundations of Genetic Algorithms XIII (2015), 150–162.

About black-box optimization and Natural Evolution Strategies

Y. Ollivier, L. Arnold, A. Auger and N. Hansen, Information-Geometric Optimization Algorithms: A Unifying Picture via Invariance Principles, J. Mach. Learn. Res., **18**(18) (2017) 1-65.

J. Martens, New insights and perspectives on the natural gradient method, J. Mach. Learn. Res., **21**(1) (2020) 5776-5851.

D. Wierstra, T. Schaul, J. Peters, and J. Schmidhuber, Natural Evolution Strategies, in: CEC 2008, IEEE World Congress on Computational Intelligence (2008) 3381-3387.

D. Wierstra, T. Schaul, T. Glasmachers, Y. Sun, J. Peters and J. Schmidhuber, Natural Evolution Strategies, J. Mach. Learn. Res., **15**(1) (2014), 949-980.

About multi-agent reinforcement learning and evolutionary games

D. Bloembergen, K. Tuyls, D. Hennes and M. Kaisers, Evolutionary Dynamics of Multi-Agent Learning: A Survey, J. Artif. Intell. Res., **53**(1) (2015), 659-697.

D. Hennes et al., Neural Replicator Dynamics: Multiagent Learning via Hedging Policy Gradient, in: AAMAS'20 Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems, (2020) 492-501.

About other applications of info-geometric approach to EGT

D. Czégel, H. Giaffar, J. B. Tenenbaum and E. Szathmáry, Bayes and Darwin: How replicator populations implement Bayesian computations, BioEssays, **44**(4) (2022).

K. Friston and P. Ao, Free energy, value, and attractors, Comput. Math. Methods Med. (2012).