Inertial Game Dynamics

R. Laraki[§] P. Mertikopoulos[*]

[§]CNRS – LAMSADE laboratory

[*] CNRS – LIG laboratory

ADGO'13 – Playa Blanca, October 15, 2013

Motivation

Main Idea: *use second order tools to derive efficient learning algorithms in games.*

The second order exponential learning dynamics (Rida's talk) have many pleasant properties, but also various limitations:

- Cannot converge to interior equilibria (not a problem in many applications, desirable in others).
- Convex programming properties not clear – no damping mechanism.
- Lack of a bona fide "heavy ball with friction" interpretation.

In this talk: use geometric ideas to derive a class of inertial (= admitting an energy function), second order dynamics for learning in games.

Approach Breakdown

The main steps of our approach will be as follows:

1. Endow the simplex with a Hessian Riemannian geometric structure.
2. Derive the equations of motion for a learner under the forcing of his unilateral gradient (taken w.r.t. the HR geometry on the simplex).
3. Derive an *isometric* embedding of the problem into an ambient Euclidean space.
4. Establish the well-posedness of the dynamics.
5. Use the system's energy function to derive the dynamics' asymptotic properties.

Notation

We will work with finite games $\mathfrak{G} \equiv \mathfrak{G}(\mathcal{N}, \mathcal{A}, u)$ consisting of:

- A finite set of players: $\mathcal{N} = \{1, \dots, N\}$.
- The players' action sets $\mathcal{A}_k = \{\alpha_{k,0}, \alpha_{k,1}, \dots\}$, $k \in \mathcal{N}$.
- The players' payoff functions $u_k \colon \mathcal{A} \equiv \prod_k \mathcal{A}_k \to \mathbb{R}$, extended multilinearly to $X \equiv \prod_k \Delta(\mathcal{A}_k)$ if players use mixed strategies $x_k \in X_k \equiv \Delta(\mathcal{A}_k)$.

Note: indices will be suppressed when possible.

Special case: if $u_{k\alpha}(x) - u_{k\beta}(x) = -\left[V(\alpha; x_{-k}) - V(\beta; x_{-k})\right]$ for some $V \colon X \to \mathbb{R}$, the game is called a potential game.

Equilibrium: we will say that $q \in X$ is a Nash equilibrium of $\mathfrak{G}$ if

$$u_{k\alpha}(q) \geq u_{k\beta}(q) \text{ for all } \alpha \in \operatorname{supp}(q_k),\ \beta \in \mathcal{A}_k,\ k \in \mathcal{N}.$$

Riemannian Metrics

A Riemannian metric on an open set $U \subseteq \mathbb{R}^m$ is a smoothly varying scalar product on U

$$g(X, Y) \equiv \langle X, Y \rangle_g = \sum_{j,k} X_j g_{jk} Y_k, \qquad X, Y \in \mathbb{R}^m,$$

where $g \equiv g(x)$ is a smooth field of positive-definite matrices on U.

Riemannian Metrics

A Riemannian metric on an open set $U \subseteq \mathbb{R}^m$ is a smoothly varying scalar product on U

$$g(X, Y) \equiv \langle X, Y \rangle_g = \sum_{j,k} X_j g_{jk} Y_k, \qquad X, Y \in \mathbb{R}^m,$$

where $g \equiv g(x)$ is a smooth field of positive-definite matrices on U.

The gradient of a scalar function $V\colon U \to \mathbb{R}$ with respect to g is defined as:

$$\operatorname{grad}_g V = g^{-1}(\partial V) \quad \text{or, in components,} \quad \left(\operatorname{grad}_g V\right)_j = \sum_k g^{-1}_{jk} \partial_k V,$$

where $\partial V = (\partial_j V)_{j=1}^n$ is the array of partial derivatives of V.

Riemannian Metrics

A Riemannian metric on an open set $U \subseteq \mathbb{R}^m$ is a smoothly varying scalar product on U

$$g(X,Y) \equiv \langle X, Y\rangle_g = \sum_{j,k} X_j g_{jk} Y_k, \qquad X, Y \in \mathbb{R}^m,$$

where $g \equiv g(x)$ is a smooth field of positive-definite matrices on U.

The gradient of a scalar function $V \colon U \to \mathbb{R}$ with respect to g is defined as:

$$\operatorname{grad}_g V = g^{-1}(\partial V) \quad \text{or, in components,} \quad \left(\operatorname{grad}_g V\right)_j = \sum_k g_{jk}^{-1} \partial_k V,$$

where $\partial V = (\partial_j V)_{j=1}^n$ is the array of partial derivatives of V.

Fundamental property of the gradient: $\frac{d}{dt} V(x(t)) = \left\langle \operatorname{grad}_g V, \dot{x} \right\rangle_g$.

More generally, the derivative of V along a vector field X on U will be:

$$\nabla_X f \equiv \langle df | X \rangle = \langle \operatorname{grad} f, X \rangle.$$

Parallel Transport

How can we differentiate a vector field along another in a Riemannian setting?

Definition

Let X, Y be vector fields on U. A connection on U will be a map $(X, Y) \mapsto \nabla_X Y$ s.t.:

1. $\nabla_{f_1 X_1 + f_2 X_2} Y = f_1 \nabla_{X_1} Y + f_2 \nabla_{X_2} Y \;\forall\; f_1, f_2 \in C^\infty(U)$.
2. $\nabla_X (a Y_1 + b Y_2) = a \nabla_X Y_1 + b \nabla_X Y_2$ for all $a, b \in \mathbb{R}$.
3. $\nabla_X (f Y) = f \cdot \nabla_X Y + \nabla_X f \cdot Y$ for all $f \in C^\infty(U)$.

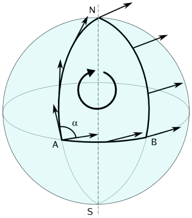

In components:

$$\left(\nabla_X Y\right)_k = \sum_i X_i \partial_i Y_k + \sum_{i,j} \Gamma_{ij}^k X_i Y_j,$$

where Γ_{ij}^k are the connection's Christoffel symbols.

Covariant Differentiation

A Riemannian metric generates the so-called Levi-Civita connection with symbols

$$\Gamma^k_{ij} = \frac{1}{2} \sum_\ell g^{-1}_{k\ell} \left(\partial_i g_{\ell j} + \partial_j g_{\ell i} - \partial_\ell g_{ij}\right)$$

This leads to the notion of covariant differentiation along a curve $x(t)$ of U:

$$\left(\nabla_{\dot{x}} X\right)_k \equiv \dot{X}_k + \sum_{i,j} \Gamma^k_{ij} X_i \dot{x}_j$$

If the field being differentiated is the velocity of $x(t)$, we obtain the acceleration of $x(t)$

$$\frac{D^2 x_k}{Dt^2} = \ddot{x}_k + \sum_{i,j} \Gamma^k_{ij} \dot{x}_i \dot{x}_j.$$

Definition

A geodesic on U is a curve $x(t)$ with zero acceleration: $\frac{D^2 x}{Dt^2} = 0$.

Hessian Riemannian Metrics

We will be interested in a specific class of Riemannian metrics on the positive orthant $\mathbb{R}^m_{>0}$ of $\mathbb{R}^m$ generated by a family of barrier functions.

Definition

Let $\theta\colon [0, +\infty) \to \mathbb{R} \cup \{+\infty\}$ be a C^∞ function such that

1. $\theta(x) < \infty$ for all $x > 0$.
2. $\lim_{x \to 0^+} \theta'(x) = -\infty$.
3. $\theta''(x) > 0$ and $\theta'''(x) < 0$ for all $x > 0$.

The Hessian Riemannian metric generated by θ on $\mathbb{R}^{n+1}_{>0}$ will be

$$g(x) = \operatorname{Hess}\left(\textstyle\sum_k \theta(x_k)\right) \quad \text{or, in components,} \quad g_{ij}(x) = \theta''(x_i)\delta_{ij}.$$

The function θ will be called the kernel of g.

Hessian Riemannian Metrics

We will be interested in a specific class of Riemannian metrics on the positive orthant $\mathbb{R}^m_{>0}$ of $\mathbb{R}^m$ generated by a family of barrier functions.

Definition

Let $\theta\colon [0,+\infty) \to \mathbb{R} \cup \{+\infty\}$ be a C^∞ function such that

1. $\theta(x) < \infty$ for all $x > 0$.
2. $\lim_{x\to 0^+} \theta'(x) = -\infty$.
3. $\theta''(x) > 0$ and $\theta'''(x) < 0$ for all $x > 0$.

The Hessian Riemannian metric generated by θ on $\mathbb{R}^{n+1}_{>0}$ will be

$$g(x) = \operatorname{Hess}\left(\textstyle\sum_k \theta(x_k)\right) \quad \text{or, in components,} \quad g_{ij}(x) = \theta''(x_i)\delta_{ij}.$$

The function θ will be called the kernel of g.

Examples

- The Shahshahani metric: $\theta(x) = x\log x \implies g_{ij}(x) = \delta_{ij}/x_j$.
- The log-barrier metric: $\theta(x) = -\log x \implies g_{ij}(x) = \delta_{ij}/x_j^2$.
- The Euclidean metric (non-example): $\theta(x) = \frac{1}{2}x^2 \implies g_{ij}(x) = \delta_{ij}$.

The Heavy Ball with Friction

The heavy ball with friction dynamics (Attouch et al.) on $\mathbb{R}^m$ are:

$$\ddot{x} = -\operatorname{grad} V - \eta \dot{x}, \qquad \text{(HBF)}$$

where $V\colon \mathbb{R}^m \to \mathbb{R}$ is a smooth potential function and $\eta > 0$ is the friction coefficient which dissipates energy.

Theorem (Alvarez 2000)

If V is convex and $\arg\min V \neq \varnothing$, (HBF) converges to a minimizer of V.

The Heavy Ball with Friction

The heavy ball with friction dynamics (Attouch et al.) on $\mathbb{R}^m$ are:

$$\ddot{x} = -\operatorname{grad} V - \eta \dot{x}, \tag{HBF}$$

where $V\colon \mathbb{R}^m \to \mathbb{R}$ is a smooth potential function and $\eta > 0$ is the friction coefficient which dissipates energy.

Theorem (Alvarez 2000)

If V is convex and $\arg\min V \neq \varnothing$, (HBF) *converges to a minimizer of V.*

We wish to apply the above method to the unit simplex Δ; in the presence of inequality constraints however, (HBF) is no longer well-posed: it exits Δ in finite time.

We will take a two-step approach:

1. Endow Δ with a Hessian Riemannian structure.
2. Derive the Riemannian analogue of (HBF).

The Heavy Ball with Friction on the Simplex

Let g be a Hessian Riemannian metric on $\mathbb{R}^{n+1}_{>0}$ with kernel θ. Then (HBF) becomes:

$$\frac{D^2 x}{Dt^2} = -\operatorname{grad}_g V - \eta \dot{x},$$

or, in components:

$$\ddot{x}_k = \frac{1}{\theta''(x_k)} u_k - \sum\nolimits_{i,j=1}^{n} \Gamma^k_{ij} \dot{x}_i \dot{x}_j - \eta \dot{x}_k,$$

with $u_k = -\partial_k V$ and $\Gamma^k_{ij} = \frac{1}{2} \frac{\theta'''(x_k)}{\theta''(x_k)} \delta_{ijk}$.

The Heavy Ball with Friction on the Simplex

Let g be a Hessian Riemannian metric on $\mathbb{R}^{n+1}_{>0}$ with kernel θ. Then (HBF) becomes:

$$\frac{D^2 x}{Dt^2} = -\operatorname{grad}_g V - \eta \dot{x},$$

or, in components:

$$\ddot{x}_k = \frac{1}{\theta''(x_k)} u_k - \sum\nolimits_{i,j=1}^{n} \Gamma^k_{ij} \dot{x}_i \dot{x}_j - \eta \dot{x}_k,$$

with $u_k = -\partial_k V$ and $\Gamma^k_{ij} = \frac{1}{2} \frac{\theta'''(x_k)}{\theta''(x_k)} \delta_{ijk}$.

Using d'Alembert's principle to project on the simplex, we obtain the inertial dynamics:

$$\ddot{x}_k = \underbrace{\frac{1}{\theta''_k} \left[u_k - \sum\nolimits_\ell \left(\Theta''_h / \theta''_\ell \right) u_\ell \right]}_{\text{Driving force}} - \underbrace{\frac{1}{2} \frac{1}{\theta''_k} \left[\theta'''_k \dot{x}_k^2 - \sum\nolimits_\ell \left(\Theta''_h / \theta''_\ell \right) \theta'''_\ell \dot{x}_\ell^2 \right]}_{\text{Constraint force}} - \underbrace{\eta \dot{x}_k}_{\text{Friction}} \quad \text{(ID)}$$

where $\theta''_k = \theta''(x_k)$ and Θ''_h is the harmonic mean $\Theta''_h = \left(\sum_\ell 1/\theta''_\ell \right)^{-1}$.

Inertial Game Dynamics

Tensoring over players, we obtain the inertial game dynamics:

$$\ddot{x}_{k\alpha} = \frac{1}{\theta''_{k\alpha}} \left[u_{k\alpha} - \sum_{\beta} \left(\Theta''_{k,h} / \theta''_{k\beta} \right) u_{k\beta} \right] - \frac{1}{2} \frac{1}{\theta''_{k\alpha}} \left[\theta'''_{k\alpha} \dot{x}^2_{k\alpha} - \sum_{\ell} \left(\Theta''_{k,h} / \theta''_{k\beta} \right) \theta'''_{k\beta} \dot{x}^2_{k\beta} \right] - \eta \dot{x}_{k\alpha}, \tag{IGD}$$

where the players' payoffs $u_{k\alpha} = \frac{\partial u_k}{\partial x_{k\alpha}}$ are viewed as unilateral gradients.

Inertial Game Dynamics

Tensoring over players, we obtain the inertial game dynamics:

$$\begin{aligned}\ddot{x}_{k\alpha} &= \frac{1}{\theta''_{k\alpha}}\left[u_{k\alpha} - \sum\nolimits_\beta \left(\Theta''_{k,h}/\theta''_{k\beta}\right) u_{k\beta}\right] \\ &\quad - \frac{1}{2}\frac{1}{\theta''_{k\alpha}}\left[\theta'''_{k\alpha}\dot{x}^2_{k\alpha} - \sum\nolimits_\ell \left(\Theta''_{k,h}/\theta''_{k\beta}\right)\theta'''_{k\beta}\dot{x}^2_{k\beta}\right] - \eta\dot{x}_{k\alpha},\end{aligned} \tag{IGD}$$

where the players' payoffs $u_{k\alpha} = \frac{\partial u_k}{\partial x_{k\alpha}}$ are viewed as unilateral gradients.

Examples

1. The Gibbs kernel $\theta(x) = x\log x$ generates the inertial replicator dynamics:

$$\ddot{x}_{k\alpha} = x_{k\alpha}\left(u_{k\alpha} - \sum\nolimits_\beta x_{k\beta}u_{k\beta}\right) + \frac{1}{2}x_{k\alpha}\left(\dot{x}^2_{k\alpha}/x^2_{k\alpha} - \sum\nolimits_\beta \dot{x}^2_{k\beta}/x_{k\beta}\right) - \eta\dot{x}_{k\alpha}. \tag{I-RD}$$

2. The Burg kernel $\theta(x) = -\log x$ generates the inertial log-barrier dynamics:

$$\ddot{x}_{k\alpha} = x^2_{k\alpha}\left(u_{k\alpha} - r_k^{-2}\sum\nolimits_\beta x^2_{k\beta}u_{k\beta}\right) + x^2_{k\alpha}\left(\dot{x}^2_{k\alpha}/x^3_{k\alpha} - r_k^{-2}\sum\nolimits_\beta \dot{x}^2_{k\beta}/x_{k\beta}\right) - \eta\dot{x}_{k\alpha}, \tag{I-LD}$$

where $r_k^2 = \sum_\beta x^2_{k\beta}$.

Energy, Damping and Convergence

For a single player, the Riemannian structure on Δ gives rise to the energy functional:

$$E(x,v) = \tfrac{1}{2}\langle v, v\rangle + V(x)$$

Under the inertial dynamics, energy is dissipated:

$$\dot{E} = \left\langle \frac{D^2 x}{Dt^2}, \dot{x} \right\rangle + \langle \operatorname{grad} V, \dot{x}\rangle = \langle -\operatorname{grad} V - \eta\dot{x}, \dot{x}\rangle + \langle \operatorname{grad} V, \dot{x}\rangle = -\eta\|\dot{x}\|^2 \leq 0$$

Energy, Damping and Convergence

For a single player, the Riemannian structure on Δ gives rise to the energy functional:

$$E(x, v) = \tfrac{1}{2}\langle v, v\rangle + V(x)$$

Under the inertial dynamics, energy is dissipated:

$$\dot{E} = \left\langle \frac{D^2 x}{Dt^2}, \dot{x} \right\rangle + \langle \operatorname{grad} V, \dot{x}\rangle = \langle -\operatorname{grad} V - \eta\dot{x}, \dot{x}\rangle + \langle \operatorname{grad} V, \dot{x}\rangle = -\eta\|\dot{x}\|^2 \leq 0$$

As a result, inertial trajectories that exist for all time eventually slow down:

Proposition

If $x(t)$ exists for all $t \geq 0$, then $\lim_{t\to\infty} \dot{x}(t) = 0$.

Energy, Damping and Convergence

For a single player, the Riemannian structure on Δ gives rise to the energy functional:

$$E(x,v) = \tfrac{1}{2}\langle v,v\rangle + V(x)$$

Under the inertial dynamics, energy is dissipated:

$$\dot{E} = \left\langle \frac{D^2 x}{Dt^2}, \dot{x} \right\rangle + \langle \operatorname{grad} V, \dot{x}\rangle = \langle -\operatorname{grad} V - \eta\dot{x}, \dot{x}\rangle + \langle \operatorname{grad} V, \dot{x}\rangle = -\eta\|\dot{x}\|^2 \leq 0$$

As a result, inertial trajectories that exist for all time eventually slow down:

Proposition

If $x(t)$ exists for all $t \geq 0$, then $\lim_{t\to\infty} \dot{x}(t) = 0$.

Theorem

Assume that the dynamics (ID) *are well-posed, and let q be a local minimizer of V with $\operatorname{Hess}(V) \succ 0$ at q. If $x(0)$ is sufficiently close to q and the system's initial kinetic energy $K(0) = \frac{1}{2}\|\dot{x}(0)\|^2$ is low enough, then $\lim_{t\to\infty} x(t) = q$.*

The Folk Theorem of Evolutionary Game Theory

First order gradient descent w.r.t. the Shahshahani metric $g_{ij}(x) = \delta_{ij}/x_j$ leads to the replicator equation:

$$\dot{x}_{k\alpha} = x_{k\alpha}\left[u_{k\alpha} - \sum_\beta x_{k\beta}u_{k\beta}\right] \tag{RD$_1$}$$

The most well known stability and convergence result is the folk theorem of evolutionary game theory which states that (RD$_1$) has the following properties:

I. A state is stationary iff it is a restricted equilibrium – i.e.$u_{k\alpha}(q) = u_{k\beta}(q)$ if $\alpha, \beta \in \operatorname{supp}(q_k)$.
II. If an interior solution orbit converges, its limit is Nash.
III. If a point is Lyapunov stable, then it is also Nash.
IV. A point is asymptotically stable if and only if it is a strict equilibrium.

An Inertial Folk Theorem

In our inertial setting, we have the following folk-type theorem:

Theorem

Assume that the inertial dynamics (IGD) *are well-posed, and let* $x(t)$ *be a solution orbit of* (IGD) *for* $\eta_k \geq 0$. *Then:*

I. $x(t) = q$ *for all* $t \geq 0$ *if and only if* q *is a restricted equilibrium.*

II. *If* $x(t)$ *is interior and* $\lim_{t\to\infty} x(t) = q$*, then* q *is a restricted equilibrium of* $\mathfrak{G}$.

III. *If every neighborhood* U *of* q *in* X *admits an interior orbit* $x_U(t)$ *such that* $x_U(t) \in U$ *for all* $t \geq 0$*, then* q *is a restricted equilibrium of* $\mathfrak{G}$.

IV. *If* q *is a strict equilibrium of* $\mathfrak{G}$ *and* $x(t)$ *starts close enough to* q *with sufficiently low speed* $\|\dot{x}(0)\|$*, then* $x(t)$ *remains close to* q *for all* $t \geq 0$ *and* $\lim_{t\to\infty} x(t) = q$.

An Isometric Embedding into Euclidean Space

The above results all rely on the inertial dynamics being well-posed – not obvious! We will study this by embedding the problem isometrically in an ambient Euclidean space.

Proposition (Nash embedding)

Let $\xi_\alpha = \phi(x_\alpha)$ with $\phi'(x) = \sqrt{\theta''(x)}$, and set

$$S = \{(\phi(x_0), \dots, \phi(x_n)) : x \in \operatorname{rel\,int}(\Delta)\}.$$

Then S with the ambient metric of $\mathbb{R}^n$ is isomorphic to $\operatorname{rel\,int}(\Delta)$ with the Hessian Riemannian metric generated by θ.

Examples

1. The open unit simplex $\Delta \subseteq \mathbb{R}^{n+1}$ with the Shahshahani metric $g_{ij} = \delta_{ij}/x_j$ is isometric to an open orthant of the radius-2 sphere in $\mathbb{R}^{n+1}$ (Akin, 1979).
2. The open unit simplex $\Delta \subseteq \mathbb{R}^{n+1}$ with the log-barrier metric $g_{ij} = \delta_{ij}/x_j^2$ is isometric to the closed hypersurface $S = \{\xi \in \mathbb{R}^{n+1} : \xi_\alpha < 0 \text{ and } \sum_\beta e^{\xi_\beta} = 1\}$.

Well-posedness of the Inertial Dynamics

In the Euclidean variables $\xi = \phi(x)$, the inertial dynamics become:

$$\ddot{\xi}_\alpha = \frac{1}{\sqrt{\theta''_\alpha}} \left(u_\alpha - \sum_\beta \left(\Theta''_h / \theta''_\beta\right) u_\beta \right) + \frac{1}{2} \frac{1}{\sqrt{\theta''_\alpha}} \sum_\beta \Theta''_h \theta'''_\beta / (\theta''_\beta)^2 \dot{\xi}_\beta^2 - \eta \dot{\xi}_\alpha.$$

By the Euclidean isometry property, this is just Newton's ordinary second law of motion for particles constrained to move on the hypersurface

$$S = \{\xi \in \mathbb{R}^{n+1} : \sum_\beta \phi^{-1}(\xi_\beta) = 1\}.$$

Theorem

The dynamics (ID) *are well-posed if and only if S is a closed hypersurface of $\mathbb{R}^{n+1}$.*

Proof technical and hard, but intuition straightforward: if S is bounded in some direction, then orbits can escape from that part of S in finite time.

Examples

Nash embedding for the Shahshahani simplex: $\theta(x) = x \log x$, $\xi = 2\sqrt{x}$.

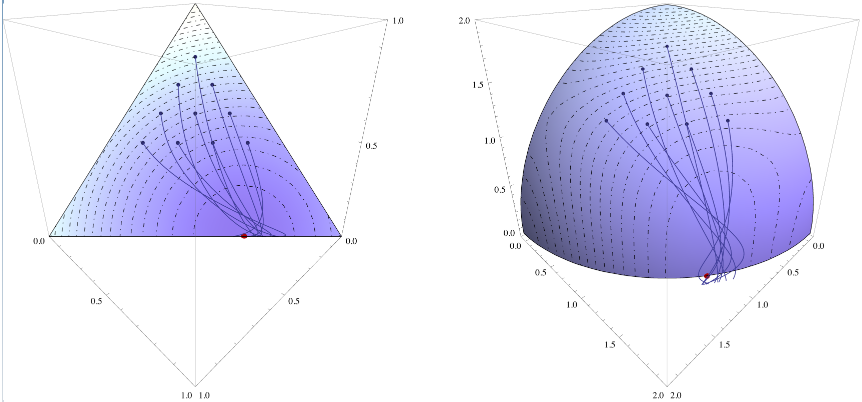

The dynamics escape in finite time.

Examples

Nash embedding for the Burg simplex: $\theta(x) = -\log x$, $\xi = \log x$.

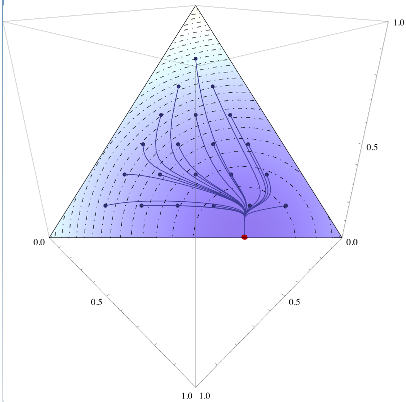

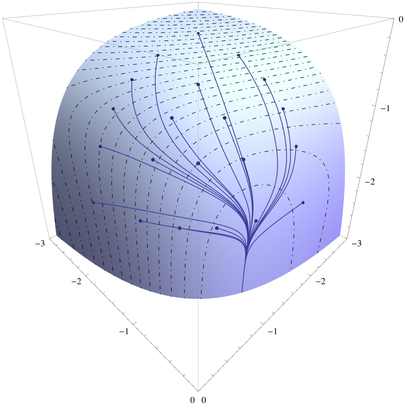

The dynamics are well-posed.

Future Directions

Some open problems for the coffee break:

- What do the dynamics look like for more general domains?
- When are they well posed?
  Conjecture: if the interior of the feasible set can be mapped isometrically to a closed submanifold of some ambient real space.
- What are the dynamics' global convergence properties for special classes of functions (convex, analytic, etc.)?
- …