The Geometry of Adversarial Learning

Mirror Flows, Poincaré Recurrence, and No-Regret Learning in Continuous Games

Davide Legacci* **Panayotis Mertikopoulos*** **Bary S. R. Pradelski***†

*Univ. Grenoble Alpes, CNRS, Inria, Grenoble INP, LIG †CNRS, Maison Française d'Oxford

Motivation: Competition Beyond Min-Max?

Potential and *min-max* games anchor two ends of a behavioral spectrum:

- *Cooperation* among players and *convergent* learning dynamics;
- *Competition* among players and *cycling* learning dynamics.

*We argue: min-max **is not** the right model for adversarial learning/competition.*

- Min-max ∩ Potential $\neq \emptyset$
- What about > 2 agents?

What is a structurally sound notion of multi-agent competition, and how should adaptive agents learn in such environments?

Hessian Flows and Poincaré Recurrence

*« In a **Poincaré-recurrent** dynamical system trajectories return **arbitrarily close to their starting points** infinitely often. »*

- ***Poincaré's + Liouville's theorems:*** $\begin{cases} v^\sharp \text{ is divergence-free} \\ \operatorname{ri}\mathcal{X} \text{ has finite volume} \end{cases} \implies \dot{x} = v^\sharp(x)$ ***is recurrent***

FTRL+: Last-Iterate Convergence and Optimal Regret

$$y_{i,n+1/2} = y_{i,n} + \hat{v}_{i,n} \qquad y_{i,n+1} = y_{i,n} + v_i(x_{n+1/2})$$
$$x_{i,n+1/2} = Q_i(\eta_{i,n+1/2}\, y_{i,n+1/2}) \qquad x_{i,n+1} = Q_i(\eta_{i,n+1}\, y_{i,n+1})$$

- ***Extrapolation step:*** Encompasses several known algorithms (DEX, Optimistic DA, …)
- ***Adaptive learn-rate:*** Environment-agnostic

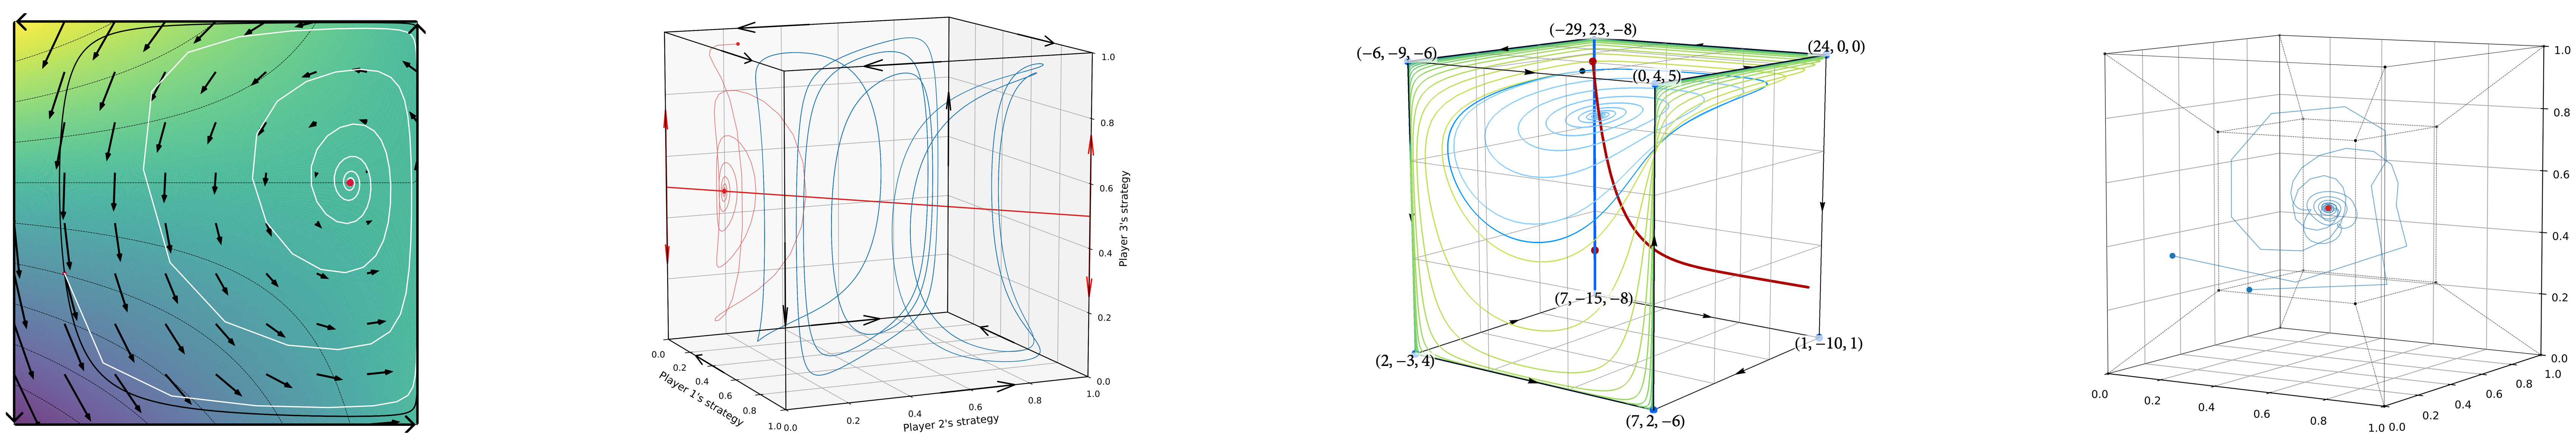

Figure: Long-run behavior of entropic mirror flow (***recurrent***) and FTRL+ dynamics (***convergent***) in 2-, 3-, and 4-player ***incompressible*** games. ***Red:*** Nash equilibria.

Continuous Games

- ***Players:*** $i = 1, \dots, N$
- ***Strategies:*** compact convex $\mathcal{X}_i \subset \mathbb{R}^{n_i}$
- ***Payoffs:*** $u_i \colon \mathcal{X} = \prod_{i \in \mathcal{N}} \mathcal{X}_i \to \mathbb{R}$
- ***Gradient field:***
$$v_i(x) = \nabla_i u_i(x) \in \mathbb{R}^{n_i}$$

The Mirror Descent Family

*« Replace Euclidean gradient steps with **mirror / Bregman** ones »*

- K-strongly convex $h \colon \mathcal{X} \to \mathbb{R}$
- Prox-mapping
$$P_x(v) = \arg\min_{q \in \mathcal{X}} \{\langle v, x - q\rangle + D(q, x)\}$$
- ***Mirror Descent***
$$x_{i,n+1} = P_{x_{i,n}}(\gamma_{i,n} v_{i,n})$$

- Choice map (soft-max)
$$Q(y) = \arg\max_{x \in \mathcal{X}} \{\langle y, x\rangle - h(x)\}$$
- ***Follow-The-Regularized-Leader***
$$y_{i,n+1} = y_{i,n} + \gamma_{i,n} v_i(x_n)$$
$$x_{i,n+1} = Q_i(y_{i,n+1})$$

Mirror Flows

- Mirror flow
$$\dot{y}_i(t) = v_i(x(t))$$
$$x_i(t) = Q_i(y_i(t))$$
- Induced ***Hessian flow***
$$\dot{x}_i = v_i^\sharp(x) := \Pi_{x_i} \operatorname{Hess} h_i^{-1} v_i(x)$$
- Riemannian metric $g_i := \operatorname{Hess} h_i$

Take-Away: Characterize and Stabilize Multi-Agent Cycles

We propose ***incompressible games*** to model continuous ***multi-agent competition***.

1. Continuous games admit a ***unique decomposition*** into potential and incompressible components $\implies$ Incompressible ∩ Potential $= \emptyset$.
2. In incompressible games, mirror flows are ***Poincaré recurrent***.
3. Optimistic dynamics recover ***last-iterate convergence*** and ***$\mathcal{O}(1)$ regret***.

Incompressible Games

*« Hessian-Riemannian **divergence** $= 0$ »*

$$0 \equiv \operatorname{div} v^\sharp(x) = \delta v(x)$$
$$= \sum_{i=1}^{N} \frac{1}{\sqrt{\det g_i(x_i)}} \sum_{\alpha_i=1}^{n_i} \frac{\partial}{\partial x_{i\alpha_i}} \left(\sqrt{\det g_i(x_i)}\, v^\sharp_{i\alpha_i} \right)$$

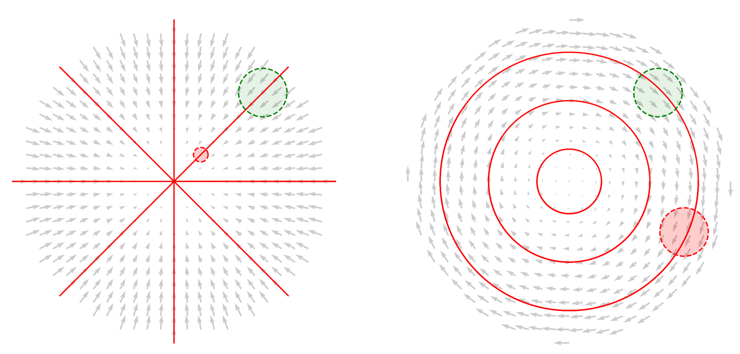

Competition Beyond Min-Max

*« Local deviations by one player induce **transverse responses** by the others »*

$$\sum_{i \in \mathcal{N}} m_i \langle v_i(x), q_i - x_i \rangle = 0$$

 # $\forall x$; $m_i > 0$; $q \in \operatorname{ri}\mathcal{X}$

- Encompasses min-max games
- ***3-player example***

*« Two players **compete** for users engagement; a third player **moderates**. »*

$$u_1(x) = (x_3 - 1/2)\, x_1$$
$$u_2(x) = (1/2 - x_3)\, x_2$$
$$u_3(x) = (x_2 - x_1)\, x_3$$

Hodge Decomposition

- On a ***compact*** manifold,
$$v = d\phi + v_{\text{inc}} \quad \text{with} \quad \delta v_{\text{inc}} = 0$$
- ***Obstacle***: $\operatorname{ri}\mathcal{X}$ *is not compact!*

We propose:
Novel compactification procedure via ***isometric smoothing map***.

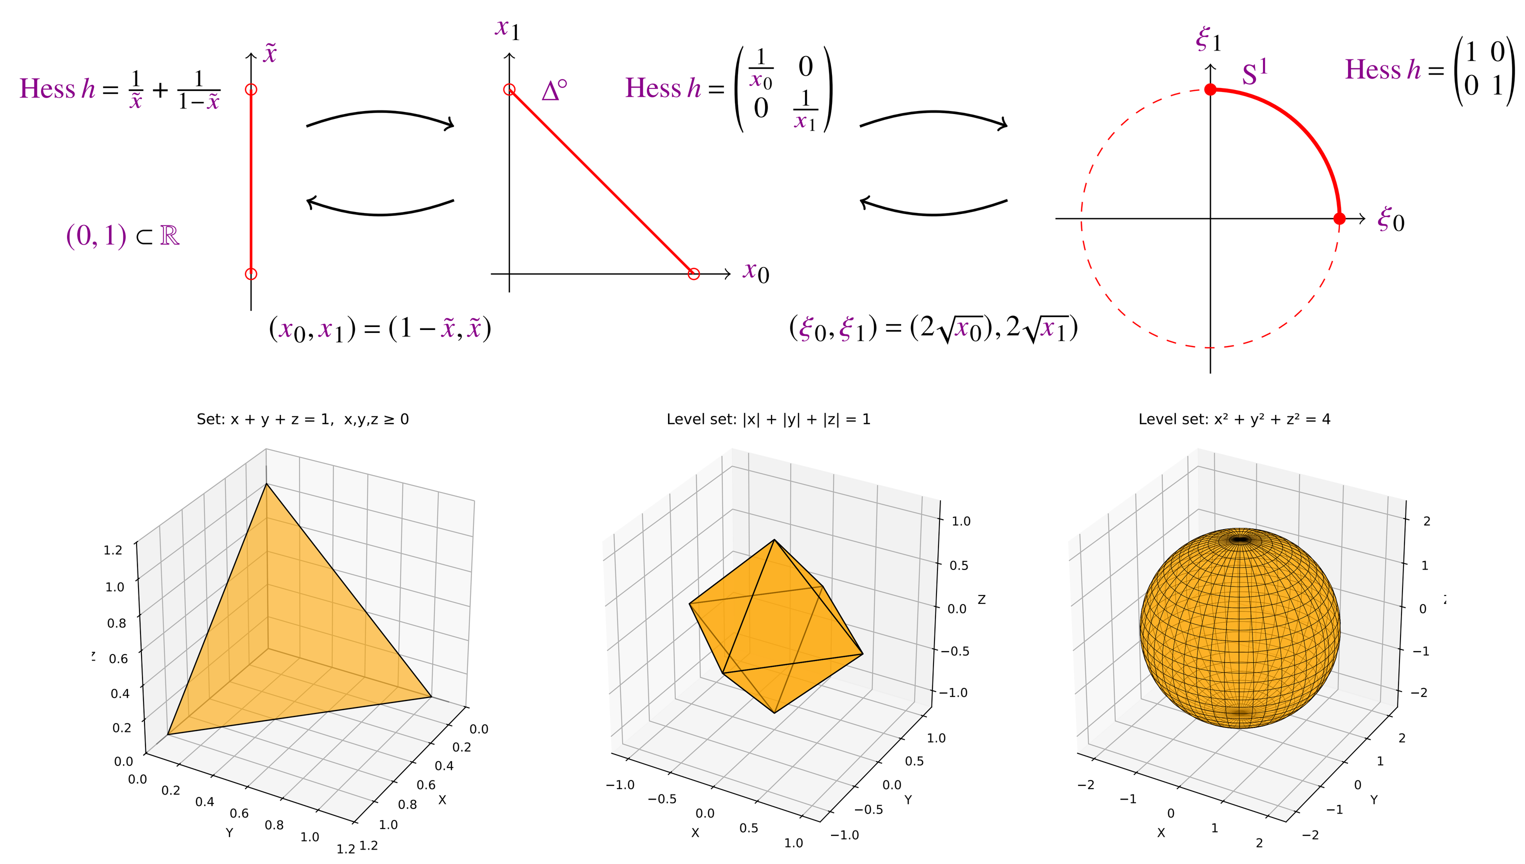

Our Contributions

Thm. — Potential-Incompressible Hodge Decomposition

The gradient field v of any continuous game can be ***uniquely decomposed*** as $v = d\phi + v_{\text{inc}}$, where $\phi \colon \operatorname{ri}\mathcal{X} \to \mathbb{R}$ is a smooth function and $\delta v_{inc} = 0$.

Thm. — Mirror Flows are Poincaré Recurrent

*In incompressible games, if $\operatorname{vol} \operatorname{ri}\mathcal{X} < \infty$, then almost every **mirror flow** orbit $x(t)$ **returns arbitrarily close to its starting point** $x(0)$ infinitely many times.*

Thm. — Last-Iterate Convergence and Optimal Regret

In finite incompressible games, ***FTRL+ converges*** to the set of Nash equilibria, guaranteeing ***constant*** $\operatorname{Reg}_i(T) = \mathcal{O}((\max h_i - \min h_i) + 1/\kappa_i)$.