Master Thesis Proposal

Non-Euclidean Geometry and Stability of Player-Wise Mirror Learning in Games

Stability, Failure Modes, and Relative Monotonicity Beyond Hilbert Space

Duration: 6 months **Type:** research thesis with theorem, counterexample, and worked examples

Positioning. Mirror descent in non-Euclidean, ℓ^p, and Banach-type geometries is classical. The thesis will therefore not claim novelty by "proving mirror descent in L^p". Instead, it asks a game-specific question: *which Hilbert-space stability mechanisms for player-wise learning in games survive under non-Hilbert Banach geometry, and which fail?* The project combines three targets into one coherent thesis: concrete geometry-dependent stability in simple games, explicit failure modes of Euclidean proof arguments, and a positive relative-monotonicity principle that explains when convergence can still be proved.

Core game model. Player i chooses $x_i \in X_i \subset E_i$, where E_i is a finite-dimensional normed space such as $(\mathbb{R}^{d_i}, \|\cdot\|_{p_i})$ or a simplex with entropy geometry. Player i has differentiable cost $\ell_i(x_i, x_{-i})$. The game field is the pseudo-gradient

$$F(x) = (D_{x_1}\ell_1(x), \ldots, D_{x_n}\ell_n(x)) \in E_1^* \times \cdots \times E_n^*.$$

A Nash equilibrium of a convex differentiable game is characterized by the variational inequality

$$\sum_i \langle D_{x_i}\ell_i(x^*), x_i - x_i^* \rangle \geq 0 \qquad (x_i \in X_i).$$

Player-wise mirror learning with mirror maps ψ_i is

$$x_{i,t+1} \in \arg\min_{x_i \in X_i} \left\{ \eta \langle D_{x_i}\ell_i(x_t), x_i \rangle + D_{\psi_i}(x_i, x_{i,t}) \right\}.$$

The game enters through F: in optimization $F = \nabla f$ comes from one objective, while in a game F is assembled from different players' objectives and may have rotational, monotone, skew, or mixed structure.

Regularizer dictionary. Each player is equipped with a regularizer / mirror map ψ_i; the learning geometry is not the norm alone, but the Bregman divergence

$$D_{\psi_i}(u_i, v_i) = \psi_i(u_i) - \psi_i(v_i) - \langle \nabla\psi_i(v_i), u_i - v_i \rangle.$$

The proposal will compare the following canonical cases, plus general strongly convex regularizers when formulating the relative-monotonicity theorem.

Geometry	Regularizer ψ	Bregman divergence	Learning rule
Euclidean	$\frac{1}{2}\|x\|_2^2$	$\frac{1}{2}\|x-y\|_2^2$	projected gradient play
ℓ^p power geometry	$\frac{1}{p}\|x\|_p^p$	p-power Bregman divergence generated by the nonlinear duality map $J_p(x) = (\lvert x_k\rvert^{p-2}x_k)_k$	p-mirror play
Simplex entropy	$\sum_k x_k \log x_k$	$\sum_k x_k \log(x_k/y_k)$, the Kullback–Leibler divergence / relative entropy	multiplicative weights / exponentiated gradient
General mirror geometry	strongly convex ψ_i on X_i	D_{ψ_i}	player-wise mirror descent

Target	Question	Expected contribution
1. Geometry-dependent stability	**How does the same game behave under Hilbert, entropy, and ℓ^p geometries?**	Derive mirror-play dynamics for bilinear zero-sum and monotone games; compare conserved, dissipated, or unstable quantities; classify special low-dimensional cases such as scalar, diagonal, or 2×2 bilinear games.
2. Hilbert-proof failure modes	**Where do Euclidean Lyapunov proofs secretly use self-duality or parallelogram identities?**	Produce explicit examples showing that Euclidean gradient-play arguments do not transfer automatically to $p \neq 2$; identify missing assumptions such as geometry compatibility, relative monotonicity, or exponent matching.
3. Relative monotonicity	**Which positive condition restores a clean mirror-learning theorem?**	Formulate a Bregman-relative monotonicity condition for games and prove a one-step Lyapunov/descent theorem for player-wise mirror learning.

Main theorem targets. The thesis should aim to deliver at least one positive theorem and one negative theorem/counterexample.

1. **Positive relative-monotonicity theorem.** For a product mirror map $\psi(x) = \sum_i \psi_i(x_i)$, prove that a suitable condition of the form

$$\langle F(x) - F(y), x - y \rangle \geq \mu\big(D_\psi(x, y) + D_\psi(y, x)\big)$$

   implies a Bregman Lyapunov descent estimate for player-wise mirror learning, under explicit step-size/smoothness assumptions.

2. **Failure-mode theorem or counterexample.** Construct a simple monotone or bilinear game showing that the usual Euclidean stability argument fails for a non-Hilbert mirror geometry unless an additional compatibility condition is imposed.

3. **Geometry-dependent classification.** For a low-dimensional bilinear zero-sum game, compare $p = 2$, $p \neq 2$, and entropy/simplex geometry; identify whether the natural energy is conserved, dissipated, or not monotone.

Running examples. The primary example is a two-player zero-sum bilinear game

$$\ell_1(x, y) = \langle Ax, y \rangle, \qquad \ell_2(x, y) = -\langle Ax, y \rangle, \qquad F(x, y) = (A^* y, -Ax).$$

This is the cleanest place where games differ from optimization: the field is rotational/skew rather than the gradient of one convex objective. The thesis should compare the Euclidean mirror map $\frac{1}{2}\|x\|_2^2$, the p-power map $\frac{1}{p}\|x\|_p^p$, and entropy on the simplex. A second example should be a monotone convex game, used to test the positive relative-monotonicity theorem.

Work packages.

1. **Dictionary.** Define convex games, pseudo-gradients, Nash-as-VI, monotonicity, and player-wise mirror maps in product spaces.

2. **Geometry computations.** Compute duality maps, Bregman divergences, KL/relative entropy, and one-step mirror inequalities for Hilbert, entropy, and ℓ^p geometries.

3. **Stability analysis.** Analyze the same bilinear and monotone games under different geometries. Identify energies/invariants where possible.

4. **Failure modes.** Isolate Euclidean proof steps that use Hilbert self-duality or parallelogram identities; turn at least one into a counterexample or impossibility statement.

5. **Relative monotonicity.** State and prove the positive theorem explaining when Bregman geometry restores a clean Lyapunov proof.

6. **Optional simulations.** Produce small plots for trajectories and Bregman energies in the running games.

Expected deliverables. A successful thesis contains: (i) a concise atlas of the relevant geometries, (ii) complete derivations of the player-wise mirror dynamics in the running games, (iii) one positive theorem based on relative monotonicity, (iv) one explicit failure-mode theorem or counterexample, and (v) a clear explanation of how games differ from optimization through the pseudo-gradient. A very strong thesis additionally classifies one low-dimensional bilinear game across several geometries.

Related work and how to use it.

Reference	How it is used
Nemirovski and Yudin, *Problem Complexity and Method Efficiency in Optimization.*	Historical source for mirror descent and non-Euclidean prox methods; use to mark the classical baseline rather than claim novelty.
Beck and Teboulle, "Mirror Descent and Nonlinear Projected Subgradient Methods" (2003).	Main finite-dimensional template for the Bregman one-step inequality; translate the proof from optimization to player-wise game learning.
Bubeck, *Convex Optimization: Algorithms and Complexity.*	Use Chapter 4, especially the mirror-descent material, for a compact derivation of the standard regret/descent proof before the game-specific extension.
Facchinei and Pang, *Finite-Dimensional Variational Inequalities and Complementarity Problems.*	Use Volume I, Chapter 1 for VIs and examples, and early Chapter 2 for monotonicity; this is the bridge from Nash equilibria to operator inequalities.
Megginson, *An Introduction to Banach Space Theory.*	Use Chapter 1 for normed spaces and duals; use Chapter 5, especially the material on convexity and smoothness, for the Banach geometry behind ℓ^p examples.
Cioranescu, *Geometry of Banach Spaces, Duality Mappings and Nonlinear Problems.*	Use the chapters on duality mappings and nonlinear problems to understand p-power duality maps as the Banach analogue of gradients.
Mertikopoulos and Zhou, "Learning in Games with Continuous Action Sets and Unknown Payoff Functions" (2019).	Use for the game-learning perspective with continuous actions and regularized learning; helpful for positioning mirror/FTRL as player dynamics.
Bauschke and Combettes, *Convex Analysis and Monotone Operator Theory in Hilbert Spaces.*	Use the chapters on convex functions, subdifferentials, and monotone operators as the Hilbert reference point; useful for identifying what breaks beyond Hilbert geometry.

Six-month plan.

Month	Milestone
1	Fix notation, choose running games, and reproduce the standard mirror inequality.
2	Compute Hilbert, entropy, and ℓ^p geometries; derive the player-wise dynamics.
3	Analyze bilinear zero-sum examples and identify candidate energies/invariants.
4	Prove the relative-monotonicity theorem in the finite-dimensional product setting.
5	Build at least one failure-mode theorem/counterexample and optional simulations.
6	Polish the thesis, check assumptions, and prepare presentation material.

Scope control. The project is research-oriented but finite-dimensional. It deliberately avoids full stochastic approximation, martingale type, weak compactness, and general infinite-dimensional L^p convergence. Those topics belong in the outlook. The core deliverable is sharper: show, in concrete games, exactly how non-Hilbert geometry changes stability, exactly where Euclidean proofs fail, and exactly which relative condition restores convergence.