Master Thesis Proposal

Discrete-Time Bregman Stability of Player-Wise Mirror Learning in Games

Building on Gao–Pavel's continuous-time relative monotonicity theorem

Duration. Six months. **Type.** Research thesis with one theorem, one counterexample/failure mode, and worked low-dimensional examples.

1 Motivation and anchor paper

Mirror descent and regularized learning are classical tools for optimization and game learning. The goal of this thesis is not to rederive mirror descent in a non-Euclidean space. Instead, the thesis will build on a specific recent result of Gao and Pavel, *Continuous-Time Convergence Rates in Potential and Monotone Games*, and ask for a discrete-time analogue in player-wise game learning.

Gao and Pavel study continuous-time dual-space dynamics in N-player concave games. In their payoff-maximization convention, with pseudo-gradient $U(x) = (\nabla_{x^p} U^p(x))_{p=1}^N$, continuous-time mirror descent has the form

$$\dot{z}(t) = \gamma U(x(t)), \qquad x(t) = C_\varepsilon(z(t)), \tag{1}$$

where C_ε is the mirror map induced by a strongly convex regularizer. Their central result is that if the game is relatively strongly monotone with respect to the Bregman geometry, then the Bregman divergence to an interior Nash equilibrium decays exponentially. In particular, if h is the product regularizer and $x^\star$ is an interior Nash equilibrium, their theorem has the form

$$D_h(x^\star, x(t)) \le \exp(-\gamma\eta\varepsilon^{-1} t) D_h(x^\star, x(0)). \tag{2}$$

Their paper explicitly warns that these continuous-time rates should not be read as discrete-time rates, because several discretizations may correspond to the same ODE and step sizes become essential in discrete time. This gap is the proposed thesis problem.

Main thesis question. Can one prove a finite-dimensional, discrete-time version of the Gao–Pavel Bregman Lyapunov theorem for player-wise mirror descent in games? If not, where exactly does the Euclidean or continuous-time proof fail?

2 Game model and sign convention

The thesis will use the cost-minimization convention. Player i chooses

$$x_i \in X_i \subset E_i, \qquad X = \prod_{i=1}^N X_i,$$

where E_i is a finite-dimensional normed space, for example $\mathbb{R}^{d_i}$ with ℓ_p geometry or a simplex with entropy geometry. Player i has a differentiable cost

$$\ell_i(x_i, x_{-i}),$$

and the game field, or pseudo-gradient, is

$$F(x) = \left(\nabla_{x_i} \ell_i(x)\right)_{i=1}^N \in E_1^* \times \cdots \times E_N^*. \tag{3}$$

In this convention $F = -U$ relative to Gao–Pavel. Thus their payoff-side inequality $(U(x) - U(y))^\top (x - y) \le 0$ becomes the cost-side monotonicity inequality

$$\langle F(x) - F(y), x - y \rangle \ge 0. \tag{4}$$

A Nash equilibrium $x^\star$ of a convex differentiable game is equivalently a solution of the variational inequality

$$\langle F(x^\star), x - x^\star \rangle \ge 0 \qquad (x \in X). \tag{5}$$

Each player is equipped with a mirror map $\psi_i : X_i \to \mathbb{R} \cup \{+\infty\}$, and the product regularizer is

$$\psi(x) = \sum_{i=1}^N \psi_i(x_i).$$

The associated Bregman divergence is

$$D_\psi(u, v) = \psi(u) - \psi(v) - \langle \nabla\psi(v), u - v \rangle. \tag{6}$$

The basic player-wise mirror descent update is

$$x_{i,k+1} \in \operatorname*{arg\,min}_{x_i \in X_i} \{\alpha \langle F_i(x_k), x_i \rangle + D_{\psi_i}(x_i, x_{i,k})\}, \qquad i = 1, \dots, N. \tag{7}$$

Equivalently, in product form,

$$x_{k+1} \in \operatorname*{arg\,min}_{x \in X} \{\alpha \langle F(x_k), x \rangle + D_\psi(x, x_k)\}.$$

3 Proposed new result

Gao–Pavel's condition is a relative, geometry-dependent replacement for ordinary Euclidean strong monotonicity. In the present cost convention, the natural discrete-time analogue is the following.

Target theorem 1 (Discrete-time relative-monotonicity descent)**.** Assume that $x^\star \in X$ solves (5), that ψ is differentiable and strongly convex on X, and that F is Bregman-relatively strongly monotone at $x^\star$ in the sense that there is $\mu > 0$ such that

$$\langle F(x) - F(x^\star), x - x^\star \rangle \geq \mu\big(D_\psi(x, x^\star) + D_\psi(x^\star, x)\big) \qquad (x \in X). \tag{8}$$

Under an explicit relative smoothness or local Lipschitz assumption on F, prove that for all sufficiently small fixed stepsizes $\alpha > 0$, the iterates (7) satisfy a one-step Lyapunov estimate of the form

$$D_\psi(x^\star, x_{k+1}) \leq D_\psi(x^\star, x_k) - c_1 D_\psi(x_{k+1}, x_k) - c_2\alpha\big(D_\psi(x_k, x^\star) + D_\psi(x^\star, x_k)\big), \tag{9}$$

with constants $c_1, c_2 > 0$ depending explicitly on the relative smoothness modulus, μ, and α.

A successful proof should imply linear convergence in Bregman divergence, for example

$$D_\psi(x^\star, x_k) \leq (1-\kappa)^k D_\psi(x^\star, x_0) \tag{10}$$

for some explicit $\kappa > 0$, at least under a symmetric-Bregman comparability condition. The thesis should state clearly which assumptions are genuinely needed. In particular, it should distinguish among:

- ordinary monotonicity (4);
- Euclidean strong monotonicity;
- Bregman-relative strong monotonicity (8);
- relative smoothness or cocoercivity assumptions needed only in discrete time.

The mathematical contribution is therefore not merely to translate Gao–Pavel's ODE proof. The goal is to identify the extra discretization terms and to determine whether they can be controlled by relative smoothness and stepsize restrictions.

4 Failure mode and counterexample target

The second required contribution is a negative result showing why the above assumptions are not cosmetic.

Target counterexample 1 (Failure of naive Euclidean transfer)**.** Construct a finite-dimensional convex game and a non-Hilbert mirror geometry such that F is monotone in the ordinary Euclidean sense, but the natural candidate Lyapunov function $D_\psi(x^\star, x_k)$ is not necessarily decreasing along (7). The example should isolate the missing compatibility assumption: either relative monotonicity, exponent matching, relative smoothness, or a Hilbert-space identity such as self-duality or the parallelogram law.

A natural test class is the two-player zero-sum bilinear game

$$\ell_1(x,y) = \langle Ax, y \rangle, \qquad \ell_2(x,y) = -\langle Ax, y \rangle, \qquad F(x,y) = (A^\top y, -Ax). \tag{11}$$

This field is skew/rotational rather than the gradient of a single convex objective. Hence it is the cleanest place to see the difference between optimization and games. In Euclidean geometry, the update resembles projected gradient play; in entropy geometry on the simplex, it becomes multiplicative weights; in p-power geometry,

$$\psi_p(x) = \frac{1}{p}\|x\|_p^p, \qquad \nabla\psi_p(x) = J_p(x) = (|x_j|^{p-2}x_j)_j,$$

the mirror map is governed by the nonlinear duality map J_p. The thesis should compute these updates explicitly in one scalar, diagonal, or 2×2 example and classify whether the natural Bregman energy is conserved, dissipated, or non-monotone.

5 Expected structure of the thesis

Chapter 1: Background. Define convex games, Nash equilibria as variational inequalities, pseudo-gradients, monotonicity, mirror maps, Bregman divergences, and the product-space notation. Include the sign conversion from Gao–Pavel's payoff convention U to the cost convention $F = -U$.

Chapter 2: The Gao–Pavel theorem. Present the continuous-time mirror descent dynamics (1), the relative monotonicity condition, and the exponential Bregman rate (2). Explain precisely why the result is continuous-time and why it does not by itself imply a discrete-time theorem.

Chapter 3: Discrete-time mirror descent in games. Derive the one-step mirror inequality for (7). Track the additional error terms that are absent in continuous time. Formulate the exact assumptions under which (9) can be proved.

Chapter 4: Positive theorem. Prove the discrete-time relative-monotonicity theorem, including explicit constants and stepsize restrictions. Whenever possible, state both a Bregman-divergence rate and a Euclidean corollary under strong convexity of ψ.

Chapter 5: Counterexample and examples. Show that ordinary monotonicity does not automatically yield Bregman Lyapunov descent in non-Hilbert geometry. Work out the bilinear game (11) in Euclidean, entropy, and p-power geometries.

6 Six-month work plan

Month	Milestone
1	Read Gao–Pavel and reproduce the continuous-time proof in the cost convention.
2	Derive the discrete-time one-step mirror inequality and identify error terms.
3	Prove a first version of the relative-monotonicity descent theorem.
4	Sharpen assumptions, constants, and stepsize restrictions; write the theorem chapter.
5	Construct the failure-mode example and compute Euclidean, entropy, and p-power cases.
6	Polish the thesis, check notation and assumptions, and prepare presentation material.

7 Expected deliverables

A successful thesis will contain:

1. a clear exposition of Gao–Pavel's continuous-time relative-monotonicity result;
2. a discrete-time player-wise mirror-descent theorem under explicit Bregman-relative assumptions;
3. at least one counterexample or failure-mode theorem showing why the assumptions are necessary;
4. worked computations for Euclidean, entropy, and ℓ_p mirror geometries in a low-dimensional game;
5. a concise explanation of how games differ from optimization through the pseudo-gradient structure.

Core references

[1] B. Gao and L. Pavel. Continuous-Time Convergence Rates in Potential and Monotone Games. SIAM Journal on Control and Optimization, 60(3):1712–1731, 2022.

[2] A. Beck and M. Teboulle. Mirror descent and nonlinear projected subgradient methods for convex optimization. *Operations Research Letters*, 31(3):167–175, 2003.

[3] A. Nemirovski and D. Yudin. *Problem Complexity and Method Efficiency in Optimization.* Wiley, 1983.

[4] P. Mertikopoulos and Z. Zhou. Learning in games with continuous action sets and unknown payoff functions. *Mathematical Programming*, 173:465–507, 2019.

[5] P. Mertikopoulos, C. Papadimitriou, and G. Piliouras. Cycles in adversarial regularized learning. *SODA*, 2018.