Novelty Assessment Report

Paper: Convergence of Regret Matching in Potential Games and Constrained Optimization

PDF URL: https://openreview.net/pdf?id=EOV1q1U23N

Venue: ICLR 2026 Conference Submission

Year: 2026

Report Generated: 2025-12-30

Abstract

Regret matching (RM)---and its modern variants---is a foundational online algorithm that has been at the heart of many AI breakthrough results in solving benchmark zero-sum games, such as poker. Yet, surprisingly little is known so far in theory about its convergence beyond two-player zero-sum games. For example, whether regret matching converges to Nash equilibria in potential games has been an open problem for two decades. Even beyond games, one could try to use RM variants for general constrained optimization problems. Recent empirical evidence suggests that they---particularly regret matching$^+$ (RM$^+$)---attain strong performance on benchmark constrained optimization problems, outperforming traditional gradient descent-type algorithms.

We show that alternating RM$^+$ converges to an ϵ-KKT point after $O_\epsilon(1/\epsilon^4)$ iterations, establishing for the first time that it is a sound and fast first-order optimizer. Our argument relates the KKT gap to the accumulated regret, two quantities that are entirely disparate in general but interact in an intriguing way in our setting, so much so that when regrets are bounded, our complexity bound improves all the way to $O_\epsilon(1/\epsilon^2)$. From a technical standpoint, while RM$^+$ does not have the usual one-step improvement property in general, we show that it does in a certain region that the algorithm will quickly reach and remain in thereafter. In sharp contrast, our second main result establishes a lower bound: RM, with or without alternation, can take an exponential number of iterations to reach a crude approximate solution even in two-player potential games. This represents the first worst-case separation between RM and RM$^+$. Our lower bound shows that convergence to coarse correlated equilibria in potential games is exponentially faster than convergence to Nash equilibria.

Core Task Landscape

This paper addresses: **Convergence of Regret Matching Algorithms in Potential Games and Constrained Optimization**

A total of **11 papers** were analyzed and organized into a taxonomy with **8 categories**.

Taxonomy Overview

The research landscape has been organized into the following main categories:

- **Theoretical Convergence Analysis**
- **Constrained Optimization via Game-Theoretic Methods**
- **Algorithm Design and Applications**

Complete Taxonomy Tree

- Convergence of Regret Matching Algorithms in Potential Games and Constrained Optimization Survey Taxonomy
- Theoretical Convergence Analysis
  - Convergence in Potential Games ★ (2 papers)
  - [0] Convergence of Regret Matching in Potential Games and Constrained Optimization (Anon et al., 2026) View paper
  - [8] Regret-based continuous-time dynamics (S. Hart, 2003) View paper
  - Convergence in General Game Settings (2 papers)
  - [10] No-Regret Learning for Stackelberg Equilibrium Computation in Newsvendor Pricing Games (Liu, 2024) View paper
  - [11] Convergent learning algorithms for unknown reward games. SIAM Journal on Control and Optimization, 51 (4), 3154-3180. DOI: 10.1137/120893501 (AC Chapman, n.d.) View paper
- Constrained Optimization via Game-Theoretic Methods
  - Non-Convex and Non-Differentiable Constraints (1 papers)
  - [1] Two-player games for efficient non-convex constrained optimization (Cotter, 2019) View paper
- Algorithm Design and Applications
  - Distributed Multi-Agent Coordination
  - Neural and Hybrid Approaches for DCOPs (2 papers)
    - [3] Neural regret-matching for distributed constraint optimization problems (Yanchen Deng, 2021) View paper
    - [7] Benchmarking hybrid algorithms for distributed constraint optimisation games (Archie C. Chapman, 2011) View paper
  - Iterative Best-Response Frameworks (1 papers)
    - [6] A unifying framework for iterative approximate best-response algorithms for distributed constraint optimization problems1 (Archie C. Chapman, 2011) View paper
  - Domain-Specific Applications
  - Dynamic Congestion and Traffic Games (1 papers)
    - [5] Distributed regret matching algorithm for dynamic congestion games with information provision (Tai-Yu Ma, 2014) View paper
  - UAV Swarm Task Allocation (1 papers)
    - [4] Task Allocation in UAV Swarm Zone Breach: A Dynamic Game-Theoretical Approach (Yanzhen Wang, 2024) View paper

◦ Imperfect Recall and Wireless Network Control (2 papers)
  - [2] Decision making under imperfect recall: Algorithms and benchmarks (E Tewolde, 2025) View paper
  - [9] Cognitive forwarding control in wireless ad-hoc networks with slow fading channels (Vesal Hakami, 2015) View paper

Narrative

Core task: convergence of regret matching algorithms in potential games and constrained optimization. The field structure divides naturally into three main branches. Theoretical Convergence Analysis examines the mathematical foundations and guarantees for regret-based learning dynamics, often focusing on potential games where convergence to equilibria can be rigorously established, as seen in foundational continuous-time formulations like Regret-based continuous-time dynamics[8]. Constrained Optimization via Game-Theoretic Methods reframes optimization problems as games, leveraging regret minimization to solve resource allocation and decision-making challenges under constraints, with applications ranging from UAV task allocation (Task Allocation in UAV[4]) to non-cooperative settings (Two-player games for efficient[1]). Algorithm Design and Applications emphasizes practical implementations, including distributed variants (Distributed regret matching algorithm[5], Neural regret-matching for distributed[3]) and domain-specific adaptations in networking (Cognitive forwarding control in[9]) and imperfect-information games (Decision making under imperfect[2]). These branches are interconnected: theoretical insights guide algorithm design, while applications motivate new convergence questions in constrained or multi-agent settings.

Particularly active lines of work explore the tension between theoretical guarantees and computational tractability. Some studies pursue unifying frameworks that connect regret matching to broader classes of learning dynamics (A unifying framework for[6]), while others benchmark hybrid approaches that blend game-theoretic methods with optimization heuristics (Benchmarking hybrid algorithms for[7]). The original paper, Convergence of Regret Matching[0], sits squarely within the Theoretical Convergence Analysis branch, specifically addressing convergence properties in potential games. Its emphasis on rigorous convergence proofs aligns closely with the continuous-time perspective of Regret-based continuous-time dynamics[8], yet it also engages with distributed and practical concerns that bridge toward works like Distributed regret matching algorithm[5]. Compared to application-driven studies such as Neural regret-matching for distributed[3], Convergence of Regret Matching[0] prioritizes foundational guarantees, contributing to the theoretical backbone that supports the field's algorithmic and applied extensions.

Related Works in Same Category

The following **1 sibling papers** share the same taxonomy leaf node with the original paper:

1. Regret-based continuous-time dynamics

Authors: S. Hart, Sergiu Hart, A. Mas-Colell, Andreu Mas-Colell | **Year/Venue**: 2003 | **URL**: View paper

Abstract

â¦ reexamine the dynamics of regret-matching from the standpoint â¦ The result for potential games is new and so we present a â¦ results of this paper), the convergence is to the set H, and not â¦

Relationship Analysis

Both papers belong to the Convergence in Potential Games category, analyzing regret matching algorithms' convergence properties in potential game settings. The original paper establishes finite-time convergence rates ($O(1/\varepsilon^4)$) for regret matching+ (RM+) to ε-KKT points in potential games and constrained optimization, while the candidate paper develops continuous-time regret-based dynamics and proves asymptotic convergence to Nash equilibria in two-person potential games. The key difference is that the original paper provides non-asymptotic complexity bounds for discrete-time RM+ with explicit iteration counts, whereas the candidate paper analyzes continuous-time differential dynamics without finite-time guarantees.

Contributions Analysis

Overall novelty summary. The paper establishes convergence guarantees for regret matching variants (particularly RM+) in constrained optimization and potential games, proving $O(1/\varepsilon^4)$ iteration complexity to ε-KKT points. Within the taxonomy, it resides in the 'Convergence in Potential Games' leaf under 'Theoretical Convergence Analysis', sharing this leaf with only one sibling paper. This represents a relatively sparse research direction focused on formal convergence proofs in potential game settings, contrasting with the more populated application-oriented branches of the taxonomy tree.

The taxonomy reveals neighboring work in 'Convergence in General Game Settings' (addressing broader game classes like Stackelberg games) and 'Constrained Optimization via Game-Theoretic Methods' (formulating optimization as games). The paper bridges these areas by treating constrained optimization through regret-based dynamics while maintaining focus on potential game structure. The taxonomy's scope notes clarify that pure game-theoretic learning without potential structure belongs elsewhere, positioning this work at the intersection of optimization theory and equilibrium computation in structured games.

Among 18 candidates examined, the contribution on RM+ convergence rates to KKT points shows one refutable candidate from 8 examined, while the connection between regret accumulation and Nash convergence shows two refutable candidates from 10 examined. The exponential separation between RM and RM+ received no examination in the limited search. These statistics suggest that while some prior work addresses related convergence questions, the specific complexity bounds and the regret-KKT gap relationship may represent less-explored territory within the examined literature.

Based on the top-18 semantic matches and citation expansion, the work appears to occupy a theoretically-focused niche with limited direct precedent in the examined candidates. The sparse population of its taxonomy leaf and the modest refutation rates suggest potential novelty, though the limited search scope means substantial related work may exist beyond these candidates, particularly in broader optimization or game theory venues not captured by this domain-specific search.

This paper presents **3 main contributions**, each analyzed against relevant prior work:

Contribution 1: Convergence rate of RM+ to KKT points in constrained optimization

Description: The authors prove that regret matching+ (RM+) converges to approximate Karush-Kuhn-Tucker points in constrained optimization problems over product of simplices with a polynomial rate. This establishes RM+ as a theoretically sound first-order optimization algorithm with complexity bounds that improve when regrets are bounded.

This contribution was assessed against **8 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. Optimization Over a Probability Simplex

URL: View paper

Brief Assessment

Optimization Over a Probability Simplex[13] focuses on a different algorithm (Cauchy-Simplex) for constrained optimization over probability simplices, not on regret matching+ (RM+) convergence to KKT points.

---

2. Wasserstein distributionally robust regret-optimal control over infinite-horizon

URL: View paper

Brief Assessment

Wasserstein distributionally robust regret-optimal[14] addresses distributionally robust control over infinite horizons using Wasserstein ambiguity sets, not regret matching convergence to KKT points in constrained optimization.

---

3. Online linear programming: Dual convergence, new algorithms, and regret bounds

URL: View paper

Brief Assessment

Online linear programming[12] focuses on dual convergence in linear programming problems with random input models, not on regret matching algorithms or their convergence to KKT points in general constrained optimization over product of simplices.

---

4. Deep learning games

URL: View paper

Prior Art Analysis

Deep learning games[18] demonstrates that regret matching+ (RM+) converges to KKT points in constrained optimization problems over product of simplices. The paper establishes a bijection between Nash equilibria in a deep learning game and KKT points of the optimization problem, and shows that RM+ can be used as a sound optimization algorithm. This prior work from 2016 predates the original paper's claims about being the first to prove RM+ convergence to KKT points with polynomial rates.

Evidence

Evidence 1 - **Rationale**: The candidate paper establishes a connection between Nash equilibria and KKT points in deep learning problems, demonstrating that RM can be used to find KKT points. This predates the original paper's claim of being the first to establish RM+ as a sound optimizer for KKT points. - **Original**: we show thatrm + converges to anϵ-kkt point afteroϵ(1/ϵ4)iterations, establishing for the first time that it is a sound and fast first-order optimizer. - **Candidate**: theorem 4 (dlg nash equilibrium) the joint action σ =(θ, {at, bt}, {qvt, dvt}) is a nash equilibrium of the dlg iff it is the joint action expansion for θ and θ is a kkt point of the dlp

Evidence 2 - **Rationale**: The candidate establishes convergence to constrained optimization solutions (which are characterized by KKT conditions) using game-theoretic methods including regret matching, providing prior theoretical foundation for RM convergence to KKT points. - **Original**: theorem 1.1.rm + converges to anϵ-kkt point of any optimization problem over a product of simplices afteroϵ(1/ϵ4)iterations. this theorem confirms thatrm + is a sound and efficient first-order optimizer - **Candidate**: ocp (one-layer constrained learning problem) add optimization constraint θ ∈θ to the olp. ocg (one-layer constrained learning game) add protagonist action constraint θ ∈θ to olg. theorem 2 (1) if (θ∗, {at, bt}) is a nash equilibrium of the ocg, then θ∗ must be a constrained global minimum of the ocp

Evidence 3 - **Rationale**: The candidate paper establishes the relationship between game utilities (related to regret) and optimization objectives (related to KKT conditions), showing how regret-based methods connect to optimization convergence. - **Original**: our argument relates the kkt gap to the accumulatedregret, two quantities that are entirely disparate in general but interact in an intriguing way in our setting - **Candidate**: lemma 3 given a fixed protagonist action θ, there exists a unique joint action for all agents σ = (θ, {at, bt}, {qvt, dvt}) where the zannis and the antagonist are playing best responses to σ. moreover ,up(σ) = -l(θ), ∇θ up(σ) = -∇l(θ)

---

5. Reducing hierarchical deep learning networks as game playing artefact using regret matching

URL: View paper

Brief Assessment

Reducing hierarchical deep learning[17] only mentions regret matching and KKT points in passing without establishing convergence rates or theoretical analysis of RM+ in constrained optimization settings.

---

6. Joint Prediction and Matching for Computing Resource Exchange Platforms

URL: View paper

Brief Assessment

Joint Prediction and Matching[15] focuses on computing resource exchange platforms with matching decisions and interior-point methods. It does not address regret matching algorithms or their convergence properties in constrained optimization or potential games.

---

7. Decision making under imperfect recall: Algorithms and benchmarks

URL: View paper

Brief Assessment

Decision making under imperfect[2] focuses on applying RM+ to imperfect-recall decision problems and benchmarking, not on proving convergence rates to KKT points in general constrained optimization.

---

8. Generalised Regret Optimal Controller Synthesis for Constrained Systems

URL: View paper

Brief Assessment

Generalised Regret Optimal Controller[16] focuses on controller synthesis for dynamic systems with disturbances, not on regret matching algorithms or their convergence to KKT points in constrained optimization.

Contribution 2: Exponential separation between RM and RM+ in potential games

Description: The authors construct two-player identical-interest games where standard regret matching (RM) requires exponentially many iterations to converge to approximate Nash equilibria, while RM+ converges polynomially fast. This is the first proven worst-case separation between these two algorithms.

This contribution was assessed against **0 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

Contribution 3: Connection between regret accumulation and convergence rate to Nash equilibria

Description: The authors establish a novel theoretical connection showing that the rate of convergence to Nash equilibria in potential games is directly governed by the regret accumulation rate. Specifically, they prove that if regret grows as T^α for α in [0,1/2], then RM+ converges to ε-KKT points in O_ε(1/ε^(2/(1-α))) iterations.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. Learning not to regret

URL: View paper

Brief Assessment

Learning not to regret[25] focuses on meta-learning regret minimizers for faster convergence on specific game distributions, not on establishing theoretical connections between regret accumulation rates and convergence to Nash equilibria in potential games.

---

2. Regret minimization in population network games: Vanishing heterogeneity and convergence to equilibria

URL: View paper

Brief Assessment

Regret minimization in population[21] focuses on smooth regret matching in population network games with heterogeneous agents, analyzing convergence to quantal response equilibria through continuity equations and moment closure techniques. The original paper examines regret matching+ (RM+) in potential games and constrained optimization, establishing direct relationships between regret growth rates (T^α) and convergence to ε-KKT points. These are fundamentally different algorithmic settings and theoretical frameworks.

---

3. Convergence of no-swap-regret dynamics in self-play

URL: View paper

Brief Assessment

Convergence of no-swap-regret dynamics[26] focuses on frequent-iterate convergence in symmetric zero-sum games under symmetric initializations, not on the general relationship between regret accumulation rates (T^α) and convergence rates to ε-KKT points in potential games or constrained optimization.

---

4. On the impossibility of convergence of mixed strategies with optimal no-regret learning

URL: View paper

Brief Assessment

On the impossibility of[27] studies impossibility of convergence to Nash equilibria in 2×2 competitive games under optimal no-regret learning, demonstrating divergence rather than convergence. This is fundamentally different from the original paper's positive convergence result showing how regret accumulation governs convergence rates in potential games.

---

5. Equilibrium tracking and convergence in dynamic games

URL: View paper

Brief Assessment

Equilibrium tracking and convergence[28] focuses on tracking Nash equilibria in time-varying games using mirror descent, not on establishing how regret accumulation rates govern convergence rates to Nash equilibria in potential games as the original paper does.

---

6. Near-optimal no-regret learning for correlated equilibria in multi-player general-sum games

URL: View paper

Brief Assessment

Near-optimal no-regret learning for[22] focuses on convergence to correlated equilibria in multi-player general-sum games using optimistic multiplicative weights update, not on the connection between regret accumulation rates and convergence to Nash equilibria in potential games or constrained optimization settings.

---

7. Regret matching+:(in) stability and fast convergence in games

URL: View paper

Prior Art Analysis

Regret matching[19] demonstrates that regret accumulation governs convergence rates to Nash equilibria in games. The paper establishes that when regret grows as T^α for α in [0,1/2], convergence to ε-KKT points occurs in O_ε(1/ε^(2/(1-α))) iterations. This directly parallels the original paper's theoretical connection, showing that the rate of convergence is parameterized by regret accumulation. Both papers prove that regret drives convergence rates to Nash equilibria in potential games, with the candidate paper explicitly stating this relationship and providing similar convergence bounds based on regret growth rates.

Evidence

Evidence 1 - **Rationale**: Both papers parameterize convergence rates as a function of accumulated regret, showing improved rates when regret is bounded, establishing the same fundamental connection between regret accumulation and convergence speed. - **Original**: our argument proceeds by parameterizing the rate of convergence ofrm + as a function of the accumulated regret, so much so that if the regret with respect to each individual simplex remains bounded, the rate is improved all the way tot -1/2 - **Candidate**: theorem 4.1. let η = d2t -1/4 and r0 = 1 . let (ft i)i∈[n] = f(zt) for t ≥ 1. for each player i, set the sequence of predictions mt i = 0 when t = 0 or restart happens at t - 1; otherwise, mt i = ft-1 i , ∀ t ≥ 1. then algorithm 1 guarantees that the individual regret regt i (ˆxi) = pt t=1 ⟨∇xiui(x...

---

8. On the convergence of no-regret learning dynamics in time-varying games

URL: View paper

Prior Art Analysis

On the convergence of[24] demonstrates prior work establishing connections between regret accumulation rates and convergence to equilibria in time-varying games. The candidate paper presents variation-dependent convergence bounds where the equilibrium gap is parameterized by regret-related variation measures, showing that sublinear regret growth implies convergence to approximate Nash equilibria. This work predates the original paper and establishes similar theoretical connections between regret properties and equilibrium convergence rates.

Evidence

Evidence 1 - **Rationale**: Both papers establish convergence rates to equilibria that depend on regret-related variation measures. The candidate's theorem 3.3 shows convergence parameterized by variation measures v(t) ϵ-ne and v(t) a, which are fundamentally connected to regret accumulation. - **Original**: theorem 1.2.suppose that the regret ofrm + on each individual simplex grows as at mostt α for someα∈[0, 1/2]. thenrm + converges to anϵ-kkt point aftero ϵ(1/ϵ 2/1-α)iterations. - **Candidate**: theorem 3.3 (detailed version in theorem a.11) . suppose that both players employ ogd with learning rate η = 1 4l in a sequence of time-varying bspps, where l := max1≤t≤t ∥a(t)∥2. then, pt t=1 eqgap(t)(z(t)) 2 = o 1 + v(t) ϵ-ne + v(t) a , where (z(t))1≤t≤t is the sequence of joint strategy profil...

Evidence 2 - **Rationale**: The candidate establishes Property 3.1 connecting dynamic regret to approximate Nash equilibria, which is a foundational element for deriving convergence rates dependent on regret accumulation, similar to the original paper's connection. - **Original**: in the special case of potential games, regret is known to drive the rate of convergence tocoarse correlated equilibria (cce); theorem 1.2 shows, for the first time, that regret can also govern the rate of convergence to nash equilibria. in particular, if convergence to cce happens at a rate oft -(1... - **Candidate**: property 3.1. suppose that z ∋z(t,⋆) = (x(t,⋆), y(t,⋆)) is an ϵ(t)-approximate nash equilibrium of the t-th game. then, for s(t) x = (x(t,⋆))1≤t≤t and s(t) y = (y(t,⋆))1≤t≤t , dreg(t) x (s(t) x) + dreg(t) y (s(t) y) ≥ -2 tx t=1 ϵ(t).

Evidence 3 - **Rationale**: Both papers establish fundamental connections between regret measures and convergence to equilibria. The candidate explicitly connects dynamic regret properties to equilibrium convergence through Property 3.1 and its generalizations, demonstrating prior theoretical work on this relationship. - **Original**: our argument relates the kkt gap to the accumulatedregret, two quantities that are entirely disparate in general but interact in an intriguing way in our setting, so much so that when regrets are bounded, our complexity bound improves all the way too ϵ(1/ϵ2). - **Candidate**: the first key observation is that dynamic regret cannot be too negative under any sequence of approximate nash equilibria (property 3.1); this was first observed by zhang et al. [93] under a sequence of exact nash equilibria. we also generalize this property by connecting it to the admission of a mi...

9. Online distributed algorithms for online noncooperative games with stochastic cost functions: high probability bound of regrets

URL: View paper

Brief Assessment

Online distributed algorithms for[23] focuses on stochastic cost functions in online noncooperative games with noisy gradients, establishing high probability bounds for dynamic regrets. The original paper establishes a novel connection between regret accumulation rates and convergence to Nash equilibria in potential games using deterministic regret matching algorithms, which is a fundamentally different theoretical contribution.

10. No-Regret Learning in Strongly Monotone Games Converges to a Nash Equilibrium

URL: View paper

Brief Assessment

No-Regret Learning in Strongly[20] focuses on Nash equilibrium convergence in strongly monotone games with continuous actions, not on potential games or constrained optimization with KKT points as in the original paper.

Appendix: Text Similarity Detection

No high-similarity text segments were detected across any compared papers.

References

- [0] Convergence of Regret Matching in Potential Games and Constrained Optimization View paper
- [1] Two-player games for efficient non-convex constrained optimization View paper
- [2] Decision making under imperfect recall: Algorithms and benchmarks View paper
- [3] Neural regret-matching for distributed constraint optimization problems View paper
- [4] Task Allocation in UAV Swarm Zone Breach: A Dynamic Game-Theoretical Approach View paper
- [5] Distributed regret matching algorithm for dynamic congestion games with information provision View paper
- [6] A unifying framework for iterative approximate best-response algorithms for distributed constraint optimization problems1 View paper
- [7] Benchmarking hybrid algorithms for distributed constraint optimisation games View paper
- [8] Regret-based continuous-time dynamics View paper
- [9] Cognitive forwarding control in wireless ad-hoc networks with slow fading channels View paper
- [10] No-Regret Learning for Stackelberg Equilibrium Computation in Newsvendor Pricing Games View paper
- [11] Convergent learning algorithms for unknown reward games. SIAM Journal on Control and Optimization, 51 (4), 3154-3180. DOI: 10.1137/120893501 View paper
- [12] Online linear programming: Dual convergence, new algorithms, and regret bounds View paper
- [13] Optimization Over a Probability Simplex View paper
- [14] Wasserstein distributionally robust regret-optimal control over infinite-horizon View paper
- [15] Joint Prediction and Matching for Computing Resource Exchange Platforms View paper
- [16] Generalised Regret Optimal Controller Synthesis for Constrained Systems View paper
- [17] Reducing hierarchical deep learning networks as game playing artefact using regret matching View paper
- [18] Deep learning games View paper
- [19] Regret matching+:(in) stability and fast convergence in games View paper
- [20] No-Regret Learning in Strongly Monotone Games Converges to a Nash Equilibrium View paper
- [21] Regret minimization in population network games: Vanishing heterogeneity and convergence to equilibria View paper
- [22] Near-optimal no-regret learning for correlated equilibria in multi-player general-sum games View paper
- [23] Online distributed algorithms for online noncooperative games with stochastic cost functions: high probability bound of regrets V iew paper
- [24] On the convergence of no-regret learning dynamics in time-varying games View paper
- [25] Learning not to regret View paper
- [26] Convergence of no-swap-regret dynamics in self-play View paper
- [27] On the impossibility of convergence of mixed strategies with optimal no-regret learning View paper
- [28] Equilibrium tracking and convergence in dynamic games View paper