STOC 2026 Paper #859 Reviews and Comments
===
Paper #859 Proximal Regret and Proximal Correlated Equilibria: A New
Tractable Solution Concept for Online Learning and Games

Review #859A
===

Overall merit
----
6. Accept (paper is among top 50% of accepted papers and I will argue for
   acceptance)

Reviewer expertise
----
3. Knowledgeable (I am fairly familiar with the area and understand the
   work reasonably well)

Paper summary

Summary of contributions
The paper studies a new notion of regret called proximal regret which sits in between external regret and swap regret. In particular, in the language of Φ-regret where Φ captures the set of endomorphisms against which the algorithm competes, prox regret is defined by setting Φ to be the set of all endomorphisms which can be written as the prox operator of a weakly convex function. The main result proves that Online Gradient Descent yields the optimal $O(\sqrt{T})$ regret rate for prox regret. As an important application of this result, the authors show that symmetric linear transformations $Ax + b$ (where A is symmetric) are a special case of prox regret, and consequently OGD minimizes regret against such endomorphisms.

The notion of prox regret naturally defines a refinement of coarse correlated equilibria called Proximal Correlated Equilibrium (PCE). When every player in a smooth game deploys OGD, the aforementioned adversarial prox regret bound implies that OGD converges to such an equilibrium with the $O(T^{-1/2})$ rate. Additionally, like external regret, the authors show that optimistic gradient descent achieves improved regret rates under self-play: when deployed by all players, optimistic gradient descent yields $O(T^{-3/4})$ prox regret for convex functions and converges to a PCE at this faster rate. Finally, the authors generalize their results to Bregman proximal regret defined using the Bregman proximal operator; here the right algorithm is Online Mirror Descent with the same distance generating function.

Strengths:
* Online Gradient Descent is one of the most important algorithms ever devised, one that is foundational to theory and ubiquitous in practice. This paper provides a fundamental improvement in our understanding of OGD by showing that it is not just another external-regret-minimizing algorithm: it minimizes the stronger and more general notion of prox regret. This notion of prox regret sits naturally between external and swap regret. Moreover, this generalization is substantial---it captures regret against the rich class of symmetric linear endomorphisms, yielding a strong regret bound that is far from obvious a priori. OGD has been shown to have special properties beyond external regret in certain applications like first-price auctions (Ahunbay and Bichler 2025; Kumar, Schneider, Sivan 2024); the current paper takes an important step towards building a theory which explains this special behavior.
* The authors show that this improvement in guarantees is robust: it extends to Bregman prox operators with OMD and continues to hold under optimistic updates for smooth convex games. Thus, these improved guarantees are not brittle and naturally generalize the known guarantees of OGD.
* The paper is extremely well written and the proofs are insightful generalizations of classic OCO bounds for external regret. It is a very nice addition to the OCO toolkit.

Weaknesses:
The paper leaves it unclear whether low prox regret is a property only enjoyed by online gradient descent, or whether other low-external-regret algorithms also achieve it. More generally, although it is nice to know that OGD has stronger regret guarantees, it is also important to understand whether or not OGD is special, i.e., understand the class of algorithms which minimize prox regret. Currently, the paper does not make an effort to analyze other popular OCO algorithms with respect to

prox regret (see comments 1 and 2).

Target Audience: Researchers interested in online learning, optimization, and game theory.

Overall Assessment: Given the importance of OGD, I believe any substantial improvement to our understanding of it---especially one as clean and insightful as low prox regret---is worthy of attention. Therefore, I am recommending accept.

Comments for authors
----
1. Please provide a discussion on the prox regret of other online learning algorithms beyond OGD. In particular, some good candidates are mean-based algorithms, FTRL (lazy projections), mismatched regularizer for Bregman divergence and prox operator etc. The current version of the paper is lacking in any guidance for algorithm design because it does not offer any comparisons.

2. Continuing from point 1, is it possible to provide a tight characterization of OGD's regret guarantees? I.e., the exact notion of regret which OGD minimizes, which could be more general than prox regret. Or alternatively, is there a tight characterization of algorithms which minimize prox regret? This seems like a difficult question in general, but perhaps it is tractable over finite actions (simplex). In particular, it would be good to know the exact location between external and swap regret inhabited by OGD's guarantees and prox regret.

3. Does your analysis say anything about dynamic regret? The original paper by Zinkevich (2003) on OGD also provides a dynamic regret bound which goes beyond external regret and allows for the benchmark to change with time (but regret depends on the amount of change), which might be worth looking into.

Typo:
* Theorem 1: "above bound hold..." should be "above bound holds..."

Review #859B
===

Overall merit
----
6. Accept (paper is among top 50% of accepted papers and I will argue for
   acceptance)

Reviewer expertise
----
3. Knowledgeable (I am fairly familiar with the area and understand the
   work reasonably well)

Paper summary

Main Contributions of the Paper:
This paper shows (Theorem 1) that the standard Canonical Gradient Descent (GD) algorithm is capable of producing vanishing \emph{proximal regret} (at essentially optimal rate) and thus converges to \emph{proximal correlated equilibria}. Proximal-regret is a special case of the notion of \phi-regret [GJ03], and sits strictly between swap regret and CE and external regret and CCE. The notion of proximal regret is related to the notion of \emph{gradient equilibrium} (which itself is related to symmetric linear swap regret) [AJT25] and the notion of \emph{semicoarse correlated equilibrium} [AB25].

The key insight leading to the main theorem (Theorem 1), establishing the $O(\sqrt{T})$ bound on the proximal regret achieved by GD is that one can follow essentially the same steps as in the classic analysis if certain terms that appear in that analysis are upper bounded in a suitable way (Lemma 1) so that these terms again telescope. A refined analysis (Theorem 2) shows that upper bounds on proximal regret imply upper bounds on symmetric linear swap regret, and GD achieves $O(\sqrt{T})$ linear swap regret.

A final set of results concerns the ‘game setting’, where players interact with each other using the same algorithm. Here the paper shows an improved individual proximal regret of $O(T^{-1/4})$ when all players employ OG (Theorem 4), and that the social regret (assuming strong convexity) is $O(1)$ (Theorem 5).

What's Best about the Paper:
- I found the result that the standard Gradient Descent algorithm actually achieves a stronger regret notion than external regret surprising and illuminating (esp. If this explains why this algorithm may find better equilibria in certain classes of games such as first-price auctions) (cf. the parallel work by Ahunbay and Bichler)

- The main theorem and its proof are almost textbook like simple

- I also enjoyed the other results, also I am less familiar with the literature here (eg the symmetric linear swap regret results and the results in the game setting)

Weaknesses:
- One thing that the paper only achieves indirectly is to prove that the concept of proximal regret actually makes a difference (for this is does seem to rely on pointing to other papers)

- For the results in the game setting it seems that there may still be some room for improvement. Can we go beyond $O(T^{-1/4})$? Can we lift the strong convexity assumption for the social regret result?

Target Audience
AGT community, learning theory community

Recommendation
For me this paper makes the cut. It's very clean, and the main result about a very fundamental algorithm.

Review #859C
===

Overall merit
----
6. Accept (paper is among top 50% of accepted papers and I will argue for acceptance)

Reviewer expertise
----
2. Some familiarity (I am moderately familiar with the area and have some understanding of the work)

Paper summary

This paper introduces the concept of proximal regret, a notion that is stronger than external regret but weaker than swap regret. It is well established that no-external-regret dynamics converge to a Coarse Correlated Equilibrium (CCE), whereas no-swap-regret dynamics converge to a Correlated Equilibrium (CE). While CCEs are computationally tractable, they can be significantly less efficient than CEs in certain games.

The authors analyze no-Φ-regret learning, a general framework where a learner's performance is compared against a sequence of actions transformed by functions from a class Φ. When Φ is the set of all transformations, this corresponds to swap regret; when Φ consists only of constant transformations, it corresponds to external regret. When all players employ no-Φ-regret algorithms, the empirical distribution of play converges to the set of Φ-equilibria.

Specifically, the authors investigate the class Φ defined by the proximal operator. Given a weakly convex function f, the transformation maps an action x to an action x' that minimizes $f(x') + \|x' - x\|^2$ (a proximal step). The authors show that when agents minimize proximal regret, the system converges to a Proximal Equilibrium, a solution concept that strictly refines the Coarse Correlated Equilibrium.

The main result establishes that Online Gradient Descent (OGD) naturally minimizes proximal regret, guaranteeing a regret bound of $O(\sqrt{T})$ for all weakly convex functions f.

Strengths:

1. Novel Solution Concept: Definition of a new class of equilibria, the Proximal Equilibrium.
2. New Learning Dynamics: Formulation of no-proximal-regret dynamics that converge to this new equilibrium class.
3. Algorithmic Guarantee: A proof demonstrating that standard OGD achieves optimal proximal regret.

Weaknesses:
1. Lack of Efficiency Analysis: The paper does not appear to discuss the efficiency (e.g., Price of Anarchy) of Proximal Equilibria compared to CCE and CE in specific game instances.

My general evaluation is that this is an elegant paper that extends the limits of learnability of equilibria very nicely using simple algorithms like OGD, and therefore, I recommend acceptance.