View Meta-Reviews

Paper ID
413

Paper Title
Tractable Local Equilibria in Non-Concave Games

Track Name
COLT2024

META-REVIEWER #1

META-REVIEW QUESTIONS

1. Verdict on this paper
Reject

2. Brief Justification.
This paper looks at some new concepts of games, and the authors prove convergence of online procedures to it.

The technical contributions are rather slim, as the techniques involved are pretty standard, but this is not (at all) an issue for such a conceptual paper.

Indeed, then main contribution to me is the introduction of a concept of "equilibria" (and Nash's paper for instance is a few pages long and a direct application of fixed point theorem). The issue is that, this is not really an equilibrium (even approximated) and the concept itself feels a bit odd to the reviewers and me.

Our suggestion would be to spend way more time on justifying these new notions so that there is no doubt that they satisfy all the wished properties (and not, as a reviewer mentioned, that it is a retro-engineering of the proof techniques). The impact of the paper would be much more important.

View Reviews

Paper ID

413

Paper Title

Tractable Local Equilibria in Non-Concave Games

Track Name

COLT2024

Reviewer #1

Questions

1. Please summarize the main claims and contributions of the paper (1-2 paragraphs).

In this paper, the authors introduce and examine the notion of an $(\epsilon,\Phi(\delta))$-local equilibrium for (constrained) non-concave games. This is a probabilistic solution concept, defined in the space of joint probability distributions over the players' strategy space, with the property that any unilateral "δ-local strategy modification" reduces the expected utility of the player in question by at most ϵ. The key in this definition is the notion of "δ-local strategy modifications", which are all functional strategy switches of the form $x_i \mapsto \phi_i(x_i)$, where ϕ_i belongs to some functional class Φ_i (the Φ in the definition above), and is at most δ-far from the identity mapping (that is, the "no-switch" map).

The authors consider two functional classes: interpolation-based, and projection-based. The first consists of all convex combinations between the identity map and a set of switching functions $\psi\in\Psi$, up to a total change of δ; the second class concerns all switches where players fix a deviation vector, and then follow that vector up to δ, projecting back to their action spaces if needed.

The authors' main contributions are as follows:
1. In the interpolation case, if Ψ consists of constant switch functions, they show that the empirical distribution of play of online gadient descent (OGD) is an $(\epsilon,\Phi(\sqrt{\epsilon}))$-local equilibrium after $\Omega(1/\epsilon^2)$ iterations.
2. In the projection case, the empirical distribution of play of OGD is an $(\epsilon,\Phi(\sqrt{\epsilon}))$-local equilibrium after $\Omega(1/\epsilon^2)$ iterations.
3. Finally, the authors are providing faster individual Φ-regret guarantees (for the functional classes in question) under optimistic gradient play with convex loss functions

(convexity is not stated in Theorem 6, but the proof of Theorem 6 is based on the guarantee of Theorem 5, which holds in the convex setting).

In terms of techniques, the authors rely on established approaches for analyzing gradient / optimistic gradient play. The main novelty is not the algorithmic analysis per se (the algorithms and techniques are fairly standard), but in tying the algorithms' guarantees to the solution concept that they introduced.

2. Describe the strengths and weaknesses of the work. Describe the significance and novelty of the contributions, relation to the prior work, clarity, and relevance to COLT.

POST-REBUTTAL
I went through the authors' rebuttal and discussed it extensively with the other reviewers, but my concerns still stand.

Regarding terminology: the crux of the matter isn't whether the proposed solution concept is a point or a distribution, but whether it captures criticality or optimality/equilibrium. In finite games, the concept of a (coarse) correlated equilibrium, approximate or otherwise, does capture the notion of an equilibrium, because of the individual concavity of the players' payoff functions. However, in the absence of individual concavity, the proposed solution concept can only capture criticality, so it is not appropriate to label it an equilibrium ("approximate", "correlated", or otherwise).

If the authors wish to emphasize that it is a "distribution" and not a "point", calling it an "approximate stationary distribution" (or "measure") would be more appropriate, as this terminology would also reflect the fact that it describes the statistics of OGD (as opposed to its limit points).

Now, to be clear, if this were just about one unluckily chosen word, I wouldn't object so strongly. However, the authors are making a point that they are proposing an equilibrium concept which is "game-theoretically meaningful, universal, and computationally tractable", so it is not a question of changing one word. I do not see how the proposed solution concept is a "game-theoretically meaningful" notion of equilibrium, so there is a conceptually significant gap here.

As for the class of phi-functions considered: in general, the relevance of phi-regret minimization depends on the class of deviation functions considered. If Φ contains all possible deviation functions, then it is a very strong concept (though, of course, one which is all but impossible to optimize); on the other hand, if Φ only contains the identity function, Φ regret minimization is a meaningless benchmark, as the comparator of each stage is the action that was actually played.

To be sure, the class $\Phi(\delta)$ considered in the paper is not "just" the identity function, but it is δ-close to it: in particular, having low $\Phi(\delta)$ regret means that, in hindsight, a player does not regret playing what they did versus something which is at most δ-far from what they actually played at each iteration. This is an important limitation, but it is not adequately discussed either, even in the authors' rebuttal.

For all these reasons, my score remains unchanged.

PRE-REBUTTAL
The paper is reasonably well written, even though notation choices like $\Phi_{\textrm{Int}}^{\mathcal{X}^i}(\delta)$ made the reading much harder (and the fact that these basic definitions are tucked away in the introduction did not help much either).

Presentation issues aside, I have the following concerns for the paper. First and foremost, the authors' solution concept is an approximate *critical point* not an approximate *equilibrium*. For example, consider the single-player (maximization) game with payoff function $u(x) = x^4$ for $x\in[-1,1]$. If I'm not mistaken, the author's definition means that a Dirac distribution at 0 would be a local equilibrium (I'm dropping here the $(\epsilon,\Phi(\sqrt{\epsilon}))$ notation to lighten notation). However, this is the worst possible action that the player could play, and *any* deviation from this point, no matter how small, would yield an improvement.

The authors are clearly aware of this, as the only "incentive guarantee" that they indicate in Table 1 is "first-order stability". However, since there is no guarantee of unilateral optimality, no matter how localized, it is misleading to label these points as approximate local equilibria (the local equilibrium notion of Ratliff and her co-authors rightly excludes cases like the above). [Incidentally, I realize that Definition 1 and the "local equilibrium" terminology may be motivated from Daskalakis et al., 2021b, 2022, but Daskalakis et al. 2023 already uses the terminology "critical points" - at least in the COLT version of the paper]

Other than this point, I do not see how the authors' solution concept captures "the convergence guarantees of GD and no-regret learning, which [they] show efficiently converge to this type of equilibrium in non-concave games with smooth utilities". For example, suppose OGD is run in the "Polar Game" of Pethick et al. (ICLR 2022, https://arxiv.org/pdf/2302.09831.pdf) or the “Forsaken Matching Pennies” example of Hsieh et al. (ICML 2021, https://arxiv.org/abs/2006.09065). For both games, the sequences generated by GD/OG converge to the same limit cycle from almost all initializations, and their time-averages converge to a point that is not stationary. In this case, how is it informative to say that "GD/OG converge to this [or any] type of equilibrium"?

Overall, I would have a more favorable view of the paper if the above limitations were made sufficiently clear to the reader and the paper's analysis provided new insights into the behavior of GD/OG in non-concave games. However, as the paper currently stands, I cannot recommend acceptance.

4. Please provide questions for authors to address during the author feedback period.

Please see above.

5. Please provide an overall score for the submission.

Weak reject

7. [Post rebuttal only] Check the box if you have read and discussed author’s feedback post rebuttal.

Yes

Reviewer #2

Questions

1. Please summarize the main claims and contributions of the paper (1-2 paragraphs).

This paper introduces a new notion of local equilibrium and show the tractability for it in non-concave games. The new local equilibrium concept the authors introduce (Definition 3) is a generalization of the (epsilon,delta)-local equilibrium introduced in [Daskalakis, '21b, Def. 3.6] and well as of a CCE for concave games. In Lemma 1, the relation between \phi(\delta)-regret and (\epsilon,\phi(\delta))-local-equilibria is shown. In particular, this gives a guarantee for convergence to a (r + \delta^2 L /2 ,\phi(\delta))-local-equilibria for non concave games, where r depends on the regret guarantees for convex losses.

The authors then showed various local strategy modifications (\phi_int and \phi_proj), for which the local (\epsilon,\phi(\delta))-equilibrium can be computed efficiently in poly(1/\epsilon) or poly(d) time (Section 4 and 5). This is achieved (a.o.) via OGD and optimistic OGD and the proofs follow relatively closely the proofs by Rakhlin and Sridharan '13. In Section 6, the authors also give an impossibility result for another natural local strategy modification \phi_Beam. In Appendix F, the authors also give hardness results (Thm. 8 and 9) for (\epsilon, \phi_int(\delta)) and (\epsilon, \phi_proj(\delta)) local equilibria.

2. Describe the strengths and weaknesses of the work. Describe the significance and novelty of the contributions, relation to the prior work, clarity, and relevance to COLT.

Strengths: - this paper is very well written and very clear.
- I like the idea of the (\epsilon,\phi(\delta))-local equilibrium and to the best of my knowledge it is new.

Weaknesses: - some of the proofs are relatively small variations of existing results (e.g. proof of Thm 3 or Thm 8,9). However, several other proofs are based on nice new ideas (e.g. the proof of Thm. 7) or an interesting new application/combination of known techniques (e.g. the proof of Thm 4), so in my opinion this is not a real weakness of the paper.

Overall: I think this paper introduces several interesting new ideas and opens up several new and interesting research directions and questions.

Minor: the notation in the proof of Thm. 3, Appendix D2, is a bit confusing. Especially the sum over f^{1:T}. For the bound on E[N_+], epsilon in (0,1) is missing. (epsilon = 1/(T sqrt(16 ln 2)) on the next page, but for this bound it is a bit confusing when you read it for the first time)

4. Please provide questions for authors to address during the author feedback period.

For which local strategy modifications \phi is the local-equilibrium concept 'meaningful' (-> question from introduction.)
More specifically: In the paragraph after Def. 5, the resemblance of the interpolation based strategy modification and CE are highlighted. Can this 'resemblance' be made rigorous? What is the interpretation of the interpolation based strategy modification if \Psi is not the set of all possible strategy modifications, but a strict subset?

%%%%%%%%% post rebuttal %%%%%%%%%%%%%%
I thank the authors for their helpful feedback. After reading and discussing with the other reviewers, I decided to vote for 'weak accept', since I can understand their objections.

5. Please provide an overall score for the submission.

Weak accept

7. [Post rebuttal only] Check the box if you have read and discussed author’s feedback post rebuttal.

Yes

Reviewer #3

Questions

1. Please summarize the main claims and contributions of the paper (1-2 paragraphs).

The paper proposes a new notion of equilibrium for non-concave games called (\epsilon, \Phi(\delta))-local equilibrium. In essence, it extends the notion of coarse-correlated equilibrium to only allow for deviates within a certain class of function \Phi(\delta). The class contains deviation functions such that || x- \phi(x) || \leq \delta. If we allow all such functions we have the notion of an (\epsilon, \delta)-local equilibrium for which Daskalakis et al proved existence and PPAD-hardness for a class of parameters. The main contribution of this paper is to show that if instead of all deviations || x- \phi(x) || \leq \delta we take one of two subclasses, here called Interpolation and Projection, then Online Gradient Descent and Optimistic Gradient both converge to an equilibrium. In other words, the paper weakens the notion of equilibrium to get positive computational results.

2. Describe the strengths and weaknesses of the work. Describe the significance and novelty of the contributions, relation to the prior work, clarity, and relevance to COLT.

My view is that the main contribution of this paper is to characterize the what is the object produced by running online gradient descent on a non-concave game. The paper is nice and well-written and gives an overview of an important literature -- I felt I learned a lot from reading. However, it feels like the paper essentially reverse engineers what should be the right definition to enable the theorem. I would find it more appealing if the authors spent more time trying to justify why their notion is natural.

My second concern is that in technical terms there is little innovation. It feels like most of the results follow from naturally extending the usual concave case using the new definition, which allow results to be ported from the concave to non-concave case more or less direclty.

4. Please provide questions for authors to address during the author feedback period.

Can you highlight the technical novelty of the paper?

Can you justify what the notion of (\epsilon, \Phi(\delta))-local equilibrium is natural? (other than making the proofs work)

5. Please provide an overall score for the submission.

Weak reject