We thank the reviewers for their thoughtful comments and feedback. Below we respond to the reviewers' concerns.

Novelty (R4&R5): The convergence rate of $O(\epsilon^{-3.5})$ sets new achievable baselines in the literature for the corresponding class of min-max games, compared to the previously best known rates in the literature; see Table 1 in the paper. In addition, the convergence analysis of the PL min-max games is new. We acknowledge that our solutions are built upon seemingly simple and intuitive techniques yet delivering better convergence results. It is worth noting that bringing these blocks together carefully and in the right order is necessary to obtain the reported rates. We also had to extend some existing techniques in non-obvious ways to establish these results. For example, Lemma A.5 (Danskin's theorem for PL functions) is not previously reported in the literature and is developed from scratch in this paper. As noted by R4, this result is novel and could be of independent interest outside the context of this paper.

Lemma A.5 (R4&R6): In the proof of Lemma A.5, on lines 429-430, we are using the Taylor expansion of $f(\cdot)$ up to the first order. Hence, we need Lipschitz smoothness of the ∇f, i.e., existence and boundedness of $\nabla^2 f$. We will further clarify this point in the updated version. Note that in this lemma (lines 431-436), we actually prove the differentiability of the function $g(\cdot)$ and its Lipschitz smoothness. The result of this lemma is not obvious due to the fact that the optimal solution set of the inner maximization is not necessarily a singleton in the PL case as opposed to the strongly concave case. We believe one source of confusion for the reviewer is that we did not explicitly state that the Taylor expansion in the proof is for function f. This led to the concern of the reviewer about the validity of Taylor expansion for function g, while this Taylor expansion is in fact for function f.

Experiments (R4&R6): We plan to expand our experimental results in multiple directions. 1) We have already performed experiments on robust training of neural networks based on a non-convex concave min-max formulation and compared our method with the state-of-the-art algorithms in this setting, namely [Madry et al 2017] and [Zhang et al 2019]. If accepted, we plan to include it in the final submission. 2) We also plan to add numerical comparisons with [37] using logistic regression (which is a convex model) (suggested by R4). 3) We used CNNs in our experiments as an example of more expressive non-convex models that are used in practice. In fact, we reported better worst case performance (higher worst class accuracy and lower variance) compared to logistic regression (R4). Note that due to the non-convexity of this setting, the algorithm of [37] would not have any convergence guarantee. However, we plan to numerically apply it to this problem as a heuristic and compare against it.

Remark 3.8 (R3): We apologize for the confusion. By stochastic gradients we mean applying SGD on the outer loop while assuming that the inner problem is solved by an oracle, i.e., it is easily solvable. For example, in our experiments the inner maximization has analytical solution when one computes the loss on each class. As you have rightfully mentioned, the convergence only applies to SGD. We reported the results for Adam in our experiments as it was more robust to the choice of step-size and thus was tuned easily. We will include this discussion alongside the results for SGD in the final version. Note that the use of SGD or Adam does not change the overall takeaways of the experiments.

Radius 1 in equations 4 & 5 (R3): Radius 1 is used for normalization. Note that the use of a finite radius in (4) and (5) guarantees that the optimum value would be bounded. But choosing radius to be 1 assures that (4) and (5) can be directly applied to the unconstrained case without any changes to the definitions, see lines 111-112.

Comparison to Dang & Lan [A] (R4): Thanks for bringing up this paper. The algorithm in [A] is developed for generalized monotone variational inequities (GMVI), which is different than our setup. To understand the difference, consider the case where $\mathcal{A}$ is singleton and Θ is unconstrained. In this case, the problem becomes an unconstrained minimization and GMVI is equivalent to $\langle \nabla_\theta f(\theta, \alpha_0) - \nabla_\theta f(\theta^*, \alpha_0), \theta - \theta^* \rangle \geq 0$. In this case, GMVI is limited in the sense that it does not cover general non-convex problems (unlike our setup). In general, neither our setup implies [A], nor [A] covers our setup. We will include a detailed discussion on the results in [A] in our revision.

PL examples (R6): Note that PL condition appears in many practical convex and non-convex problems; see lines 71-74 in the paper. In fact, it is more general than strong convexity because it allows the existence of multiple optimal solutions, e.g., if $f(\cdot)$ is strongly convex and A is a linear mapping, $f(A(\cdot))$ is PL. Thus, any practical problem where the inside max is concave, but satisfy such a form, e.g. high-dimensional linear or logistic regression, would be a PL game. In the non-convex optimization though, the examples are more specific; see [15, 17, 48, 27] for some. Example 3.1, i.e., generative adversarial imitation learning for LQR, is a great example and we definitely plan to expand it and explain it more explicitly. It is worth noting that using non-convex game formulations in learning applications is relatively new. We believe more specific examples would be found as the field is moving forward very quickly.

Minor comments: We will expand Table 1 to include more details of all the algorithms (R3). We will add more details on how we obtained the expressions in lines 449 and 530 (R4). We agree that we should have cited Nemirovski's paper on convex-concave problems. Thanks for bringing it up. We will cite it appropriately in our final version (R4).

[A] C. D. Dang and G. Lan. On the convergence properties of non-euclidean extragradient methods for variational inequalities with generalized monotone operators.