

1 We thank the reviewers for their time, effort, and helpful feedback. We will make sure to incorporate the reviewers’  
2 suggestions in the final version of the paper. We address individual feedback below. The citation keys are the same with  
3 the main paper, with one additional reference for this rebuttal.

4 **Reviewer 1:** “... a major takeaway I get is that the PGD attack seems to provide an okay approximation of robustness ...  
5 I would be interested in seeing how the other verifiers (which do not fall within the convex-relaxed framework) are  
6 compared to the PGD attacks.”

7 Reply: First, we affirm that “the PGD attack seems to provide an okay approximation of robustness compared to the  
8 layer-wise convex-relaxed methods” is an empirical observation, based on our empirical results in Table 1 and empirical  
9 results in other papers, e.g., [Tjeng et al., 2019, Xiao et al., 2019]. However, these empirical results are only obtained  
10 for small-size networks, where obtaining the exact answer is feasible via MILP.

11 Second, verifiers outside our framework, e.g., [Raghunathan et al., 2018b, Bunel et al., 2018, Singh et al., 2019b],  
12 may provide tighter upper bounds as compared to the upper bounds of the relaxed verifiers within our framework.  
13 However, these verifiers are significantly more computationally expensive than LP relaxed classifiers (which already  
14 take 22 CPU-years to benchmark in our paper), and thus a careful and comprehensive experiment should be designed to  
15 benchmark them. It is one important future work to be done in this field, but it is outside the scope of our paper.

16 After all, PGD is a more **practical** approximation of robustness for networks, compared to convex relaxation based  
17 methods. However, for a randomized smoothing classifier (which uses a base network but is not a network itself, see,  
18 e.g., Cohen et al., ICML 2019), the certified robustness provided by Cohen et al. is easy to compute and tight, and can  
19 **also** serve as a good approximation of robustness for randomized smoothing classifiers. In addition, we emphasize that  
20 PGD provides **lower bounds** on the robust error while relaxed-verifiers provide **upper bounds**. We should be cautious  
21 when comparing these two different types of bounds.

22 Please consider raising your score if you like our work.

23 **Reviewer 2:** Thank you for your positive review!

24 **Reviewer 3:** “Please change figures to high resolution.” Reply: Thanks, will do in the revised version.

25 “Equations 3 - It might be clearer to spread it on more rows. Explaining each constraint in this formulation might help  
26 readability.” Reply: Thanks, will reformat the equation and add explanations and intuitions in the revised version.

27 “Proofs and discussions of relaxation relationships in Figure 1” Reply: The edges labeled with Theorem 4.2 and  
28 Corollary 4.3 are our main theoretical contributions. In the following, we provide discussions about other relationships,  
29 which we will make clear and provide pointers to places in the appendix in the revised version.

30 “Optimal layer-wise convex relaxation -> CROWN”: trivially holds since CROWN is a greedy algorithm to solve LP  
31 relaxations (problem  $\mathcal{C}$  plus Eq. (7)), which can be included in the convex relaxation framework (Appendix D).

32 “CROWN -> Fast-Lin”: Zhang et al., NIPS 2018 proposed CROWN as a more general variant of Fast-Lin. In Fast-Lin,  
33 the linear relaxation needs to use the same slope for upper and lower bounds; in CROWN, the slope can be different. In  
34 other words, in Eq. (7) the  $\bar{a}^{(l)} = \underline{a}^{(l)}$  for Fast-Lin but this is not a requirement anymore for CROWN.

35 “LP-Relaxed Dual -> CROWN” and “LP-Relaxed Dual -> Fast-Lin”: CROWN and Fast-Lin uses one linear upper  
36 bound and one linear lower bound (Eq. (7)), so problem  $\mathcal{C}$  becomes a special case of LP-relaxed problem. Especially,  
37 line 183-188 and line 572-576 discussed the relationship between dual-LP, Fast-Lin and CROWN.

38 “DeepPoly <-> CROWN” and “DeepZ <-> Fast-Lin”: A simple comparison of these algorithms would show them  
39 to be the same. While different papers use different notations, one can straightforwardly translate between them. In  
40 (Singh et al., ICLR 2019) page 9, section 4.2, the author (same author as DeepZ and DeepPoly) commented “We note  
41 that DeepZ has *the same* precision as Fast-Lin (Weng et al., ICML 2018) and DeepPoly has *the same* precision as  
42 CROWN (Zhang et al., NIPS 2018).” Also in the experiments of DeepZ (Singh et al., NIPS 2018), the lines for DeepZ  
43 and Fast-Lin overlap (exactly the same values are obtained). Note that these algorithms have slight differences (for  
44 example, DeepPoly and DeepZ consider floating point rounding issues and have a faster implementation), but the basic  
45 algorithm computes exactly the same bounds.

46 “DeepPoly -> DeepZ”: Relationship is similar to “CROWN -> FastLin”. Line 553-565 in Appendix D briefly discussed  
47 the relaxation relationship between these algorithms.

48 “Optimal layer-wise convex relaxation -> DeepPoly”: holds because CROWN and DeepPoly use the same relaxation  
49 and algorithm to greedily solve the resulting linear programming problem.

50 “Fast-Lin <-> Neurify”: Fast-Lin and Neurify uses the same relaxation for ReLU neurons (and unlike other works, these  
51 two only deal with ReLU activation functions). This can be observed by comparing figure 3 in Neurify paper and Figure  
52 1 in the Fast-Lin paper: the selection of slope  $\underline{a}^{(l)}$  and  $\bar{a}^{(l)}$  are the same (line 513 in Appendix D).

53 **Additional references** Cohen, J.M., Rosenfeld, E. and Kolter, J.Z.. Certified adversarial robustness via randomized  
54 smoothing. ICML 2019