

1 **Authors' Response for NeurIPS Paper Number 6086.** We thank the reviewers for their constructive comments. We
2 first provide a general response common to all the reviewers and then respond to each reviewer's particular comments.

3 **Beyond MNIST and Feed-Forward Fully-Connected (FFC):** Since the submission of the paper, we have moved
4 beyond MNIST and FFC neural networks; in particular, we have performed experiments on several CNNs trained on
5 CIFAR-10 and confirmed that (a) our method can scale to larger and more complex networks*, and (b) our bounds are
6 consistently superior to the competing approaches. We will release the code necessary to reproduce these results. In its
7 current status our approach scales to non-trivial tasks such as deep RL for control or robotics [5,6,7]. We can further
8 improve the scalability of our tool by simply improving our hardware (computations were performed on a basic laptop)
9 and our software (exploiting the structure in the SDP, for example).


10 *We investigated CNNs with two convolutional layers (with three and six filters, no padding, and a stride of one)
11 followed by one linear layer. We remark that unrolling the convolutional operators in a CNN results in a (large)
12 feedforward network. This conversion implicitly handles the padding and stride hyperparameters. Furthermore, since
13 the max function is convex, we can describe the max pooling operation using quadratic constraints without additional
14 assumptions. Therefore, we can directly use LipSDP for CNNs. The CNNs we used had approximately 15,000 neurons
15 after unrolling.

16 **Reviewer 1: Software release:** We have already implemented the software necessary to reproduce our results. All of
17 the code is included in a private GitHub repository that will be made public with the camera-ready version of this paper.
18 Our software links to both interior-point method solvers (Mosek) and first-order method solvers, which are specifically
19 tailored for large-scale cone programs. We tried several large-scale solvers but found the SCS solver, developed by
20 Boyd's group, to be the most stable one.

21 **Reviewer 2: Lipschitz constant w.r.t. other norms:** In our experiments, we bounded the Lipschitz constant in ℓ_2 , then
22 derived the adversarial balls in ℓ_2 , and finally converted the ℓ_2 balls to ℓ_∞ balls using the inequality $\|x\|_p \leq n^{\frac{1}{p}-\frac{1}{q}} \|x\|_q$
23 ($x \in \mathbb{R}^n$). In Remark 1, we discussed how to convert the Lipschitz bound from ℓ_2 norm to other ℓ_p -norms. Apart from
24 this norm conversion, we believe that reformulating the problem (by, e.g., moving to the dual domain or changing
25 the objective function) will enable us to incorporate other ℓ_p norms ($p = 1, 2, \infty$) directly in LipSDP. A thorough
26 treatment of this is an important future research direction. **Parallelization:** In principle, we agree with the reviewer
27 that the splitting+parallelization scheme discussed in our paper is not exclusive to our method. However, among the
28 methods that can split the computation, we obtain more accurate bounds as our bounds are more accurate per chunk.

29 **Reviewer 3: More experiments:** (a) As suggested by the reviewer, we used the training method in [2] to
30 train an FFC NN with architecture 784-200-100-10 on MNIST for various values of λ (a hyper-parameter
31 controlling the Lipschitz constant of each layer). The figure on the right shows that LipSDP yields tighter
32 bounds than those guaranteed by [2]. (b) We trained various NNs on MNIST with different initializations.

33 The variance of the Lipschitz bounds found by our method was essentially negligible. For instance, with one hidden layer and 200 neurons,
34 the Lipschitz bounds were around 20 and the standard deviation over 6 trials was < 0.01 . **Naive lower bound:** In the special case of a purely
35 linear network or under certain positivity assumptions, the "naive lower bound", defined in [4], is equal to the Lipschitz constant. By
36 testing on randomly generated networks, it is our impression that this naive bound, introduced in [4], may be neither a lower bound
37 nor an upper bound on the Lipschitz constant. We will clarify this in the revised manuscript. **Presentation:** We have improved the nota-
38 tion, wording, technical presentation, bibliography, and figures in the revised manuscript. For instance, (a) we clarified that Φ is "the
39 concatenation of activation functions at each layer"; (b) we made the presentation of the assumptions about the activation
40 functions consistent and mentioned that [1] violates the assumption in our paper; (c) we updated the references and
41 cited [1] and [2].



48 **References:** [1]: Anil, Cem, James Lucas, and Roger Grosse. "Sorting out Lipschitz function approximation." [2]:
49 Gouk, Henry, et al. "Regularisation of neural networks by enforcing Lipschitz continuity." [3]: Dhillon, Guneet S., et al.
50 "Stochastic activation pruning for robust adversarial defense." [4] Combettes, Patrick L., and J. Pesquet. "Lipschitz
51 Certificates for Neural Network Structures Driven by Averaged Activation Operators." [5]: Lillcrap, Timothy P., et al.
52 "Continuous control with deep reinforcement learning." [6]: Shi, Guanya, et al. "Neural lander: Stable drone landing
53 control using learned dynamics." [7]: Duan, Yan, et al. "Benchmarking deep reinforcement learning for continuous
54 control."