Exercise 04

DeepPoly Branch and Bound Certification

Reliable and Trustworthy Artificial Intelligence
ETH Zurich

Problem 1 (DeepPoly Branch and Bound)**.** Consider the neural network below, taken from this week's lecture slides. We show the result of analysing the network using the DeepPoly algorithm on the ℓ_∞ region $\left\| \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} \right\|_\infty \leq 1$ i.e. ℓ_∞ ball around $\begin{bmatrix} 0 \\ 0 \end{bmatrix}$ with size 1.

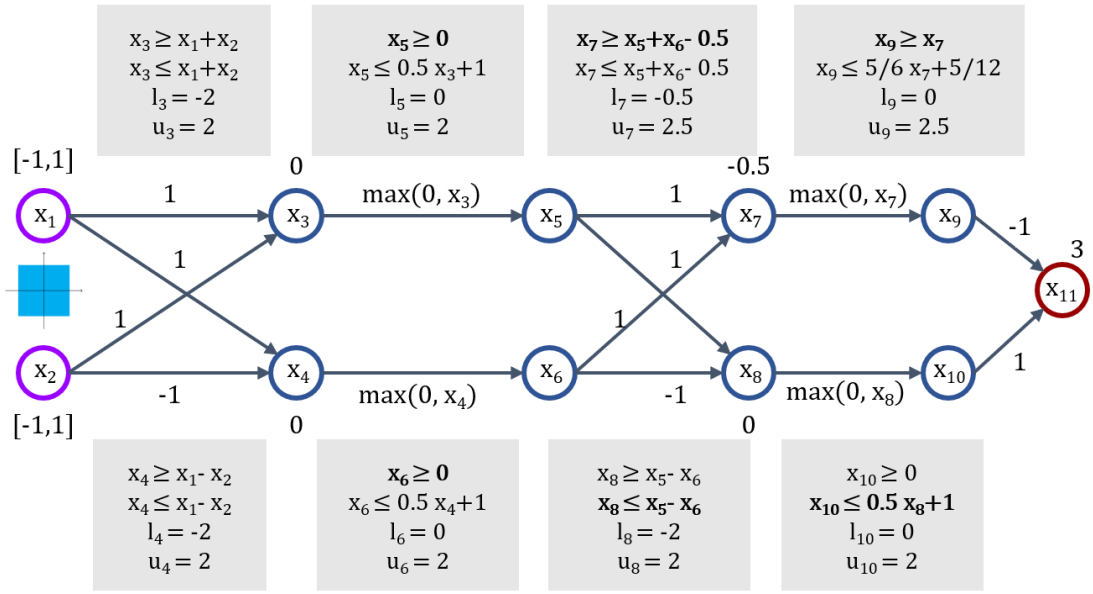

(a) Recall from the lecture, in the original DeepPoly analysis we computed the upper bound of x_{11} to be 4.5. Apply branching to the ReLU node at x_8. What upper bound for x_{11} do you obtain if you apply symbolic analysis on β (where β is the KKT variable introduced by the split at x_8, as in the lecture)? Is the resulting bound more or less precise than the original bound?

(b) The analysis you performed in (a) was done for the input region represented by an ℓ_∞ ball of size 1 around $\begin{bmatrix}0\\0\end{bmatrix}$. Without changing the intermediate neuron lower and upper bounds, use the Holder inequality to similarly compute an upper bound on x_{11} for two additional input regions — an ℓ_1 and ℓ_2 balls of size 1 around $\begin{bmatrix}0\\0\end{bmatrix}$. Is the resulting upper bound on x_{11} sound? How can you make it more precise?

(c) In (a) and (b), we applied symbolic analysis to obtain the upper bound on x_{11}. This is often infeasible in practice. Next, we find the value of β that produces the best upper bound for x_{11} using numerical optimization for the original ℓ_∞ input region. Assume, β is initialized to 1.2 (for both branches). Perform one gradient step on β with step size 0.3 on both branches. What upper bound do you obtain for x_{11}? Is the produced upper bound sound? How does it compare to the original DeepPoly bound? How does it compare to the bound obtained in (a)?