

A Appendix

A.1 Proof of Theorem 3.1

Proof. Since $P_{k\text{-ReLU},i}$ contains $\mathcal{P}_{R,i} \neq \emptyset$ and both $P_{k\text{-ReLU},i}$, $\bigcap_{i=1}^k P_{1\text{-ReLU},\mathcal{I}_i}$ contain the same interval constraints, we have that $P_{k\text{-ReLU},i} \subseteq \bigcap_{i=1}^k P_{1\text{-ReLU},\mathcal{I}_i}$. We have by monotonicity of the intersection and the convex hull,

$$\left(\bigcap_{S_i \in \mathcal{S}_i} P_{k\text{-ReLU},i} \cap S_i \right) \subseteq \left(\bigcap_{S_i \in \mathcal{S}_i} \left(\bigcap_{u \in \mathcal{I}_i} P_{1\text{-ReLU},u} \right) \cap S_i \right) \quad (9)$$

For any $S_i \in \mathcal{S}_i$, we have that either $S_i \subseteq \mathcal{C}_u^+$ or $S_i \subseteq \mathcal{C}_u^-$ for any $u \in \mathcal{I}_i$. Thus, we can replace all S_i on the right hand side of (9) with either \mathcal{C}_u^+ or \mathcal{C}_u^- such that for all $u \in \mathcal{I}_i$ both \mathcal{C}_u^+ and \mathcal{C}_u^- are used in at least one substitution and obtain by monotonicity,

$$\begin{aligned} &\subseteq \left(\bigcup_{j=1}^k \left(\left(\bigcap_{u \in \mathcal{I}_i} P_{1\text{-ReLU},u} \right) \cap \mathcal{C}_u^+ \right) \cup \left(\left(\bigcap_{u \in \mathcal{I}_i} P_{1\text{-ReLU},u} \right) \cap \mathcal{C}_u^- \right) \right) \\ &\subseteq \left(\bigcup_{u \in \mathcal{I}_i} (P_{1\text{-ReLU},u} \cap \mathcal{C}_u^+) \cup (P_{1\text{-ReLU},u} \cap \mathcal{C}_u^-) \right). \end{aligned}$$

For other i 's, it can be shown similarly that $(\bigcap_{S_i \in \mathcal{S}_i} P_{k\text{-ReLU},i} \cap S_i) \subseteq (\bigcup_{u \in \mathcal{I}_i} (P_{1\text{-ReLU},u} \cap \mathcal{C}_u^+) \cup (P_{1\text{-ReLU},u} \cap \mathcal{C}_u^-))$ holds. Since \subseteq relation holds for at least one i and \subseteq holds for other i 's, $P_{k\text{-ReLU}} \subseteq P_{1\text{-ReLU}}$ holds by the order preserving property of the intersection. \square

A.2 1-ReLU vs 2-ReLU vs 3-ReLU on the 9×200 network

We graphically show the precision and the average runtime of 1-ReLU, 2-ReLU, and 3-ReLU with no MILP and with one layer of MILP on the 9×200 MNIST network with $\epsilon = 0.015$ using the results from the evaluation section in Fig. 4.




Figure 4: Percentage of images verified and the average runtime (in seconds) of 1-ReLU, 2-ReLU, and 3-ReLU on the 9×200 MNIST network with $\epsilon = 0.015$.