Structural or analytic derivation

The cube support function $h_{[0,1]^D}(w) = \sum_j (w_j)_+$ replaces enumeration of every possible input. It gives a global activation certificate and identifies precisely the boundary at which derivative conventions matter.

Comparative cost and preferred method

The certificate takes $O(D)$ arithmetic and one pass over the weights, compared with evaluating 2^D inputs at $O(D)$ each. The pass is worst-case necessary in an unrestricted coefficient-access model: an uninspected coefficient can be changed to a sufficiently large positive value and alter the answer. This is a genuine exact complexity reduction; it proves activation behaviour, not a universal assertion about every optimizer called gradient-based.