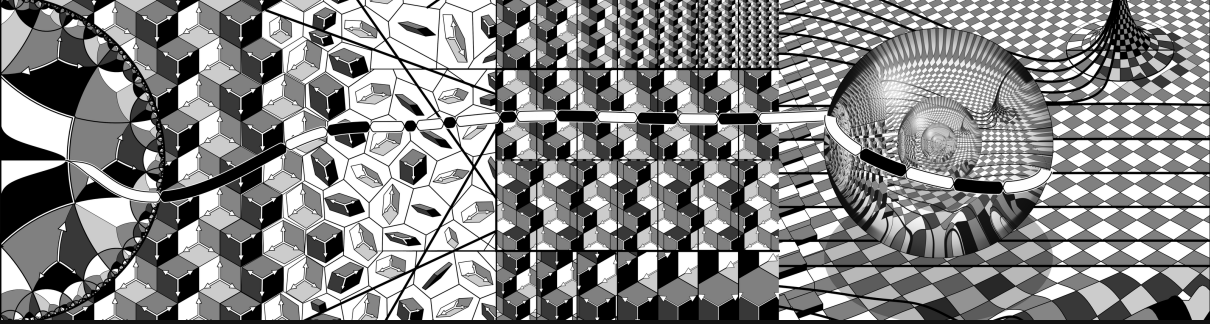


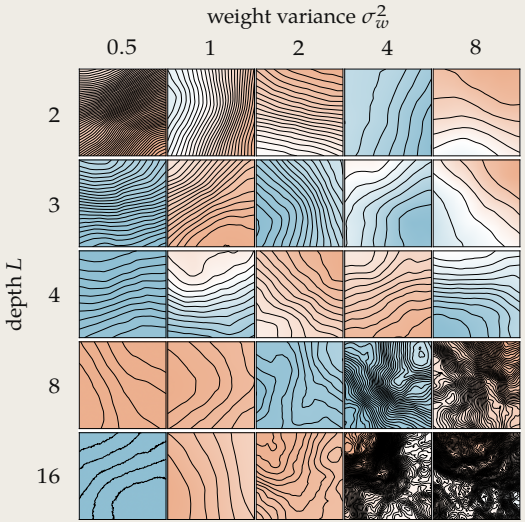
NEURIPS 2026

Tropical Boundary Complexity of Deep ReLU Binary Classifiers

Jun Li

Australian Artificial Intelligence Institute, University of Technology Sydney

How does depth shape the decision boundary?



One random ReLU network per panel.
Thin curves: level sets of the log-odds.
Thick line: the decision boundary.
Rightwards: larger weight variance.
Downwards: more layers.

What does one more layer do to the boundary, and which initialization parameters control it?

Each unit of the last hidden layer bends the boundary

At depth two the boundary is determined by two zonotopes, with one generator per hidden unit. Zhang et al. 2018; Alfarra et al. 2023

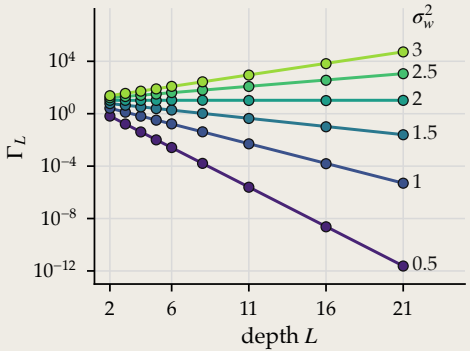
Generator energy. One term per unit j:

$$Z = \|\mathbf{G}_1\|_F^2 + \|\mathbf{G}_2\|_F^2 = \sum_j \beta_j^2 \left\|A_{j,:}\right\|_2^2 . \tag{1}$$

β_j: jump in the slope of the log-odds when unit j switches on.
$A_{j,:}$: input weights of unit j.

Its mean over random networks is proportional to the mean squared gradient of the log-odds. This is the quantity we follow as layers are added.

Result A: each layer contributes a factor $\sigma_l^2/2$ to the energy



Lines: formula. Markers: Monte Carlo, width 128.

Γ_L: mean energy at depth L, over the weights and the input.

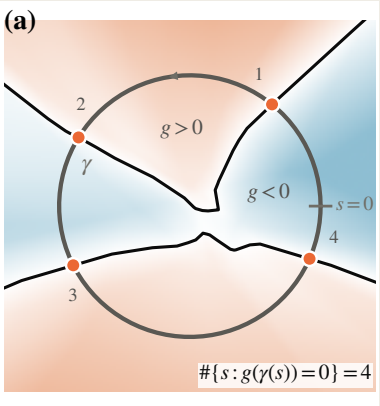
$$\Gamma_L = 8\Big(1 + \frac{1}{\pi}\Big) \prod_{l=1}^{L} \frac{\sigma_l^2}{2} \tag{2}$$

Exact at finite width. Bias variances do not enter.

He initialization, $\sigma_l^2 = 2$: constant in depth. Below 2: decays. Above 2: grows.

One network stays close to Γ_L when the width is large compared with the depth.

Result B: the crossing rate is set by the ratio of energy to output variance



A great circle γ in the input space, the boundary $g = 0$, four crossings.

Count how often the boundary crosses a great circle in the input space. The rate per unit length is

$$\rho_L = \kappa(p_0)\sqrt{\frac{\Gamma_L}{q^{(L)}}} \tag{3}$$

$\kappa(p_0)$ depends on the input dimension only.

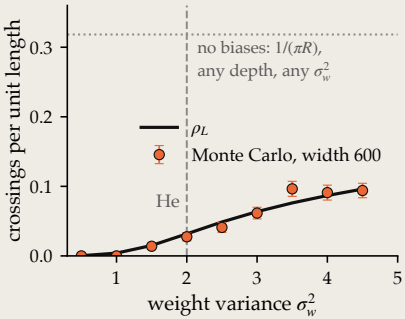
Γ_L: how fast the log-odds changes. $q^{(L)}$: the output variance Poole et al. 2016, how far the log-odds is from zero.

From Rice's formula: given the lower layers, g and its derivative along the circle are Gaussian.

Result B: biases lower the crossing rate, and the weight variance sets by how much

With biases, $q^{(L)}$ grows while Γ_L does not, so the ratio drops with depth and the boundary is crossed less often. The drop is largest at small weight variance, and the prediction matches simulation.

Baseline without biases: the ratio is constant in depth, $\rho_L = 1/(\pi R)$, two crossings per great circle as for one hyperplane.



$L = 8$, $\sigma_b^2 = 0.5$. Points: Monte Carlo, width 600.

Code and data
`github.com/junjy007/tropical-boundary-complexity`