

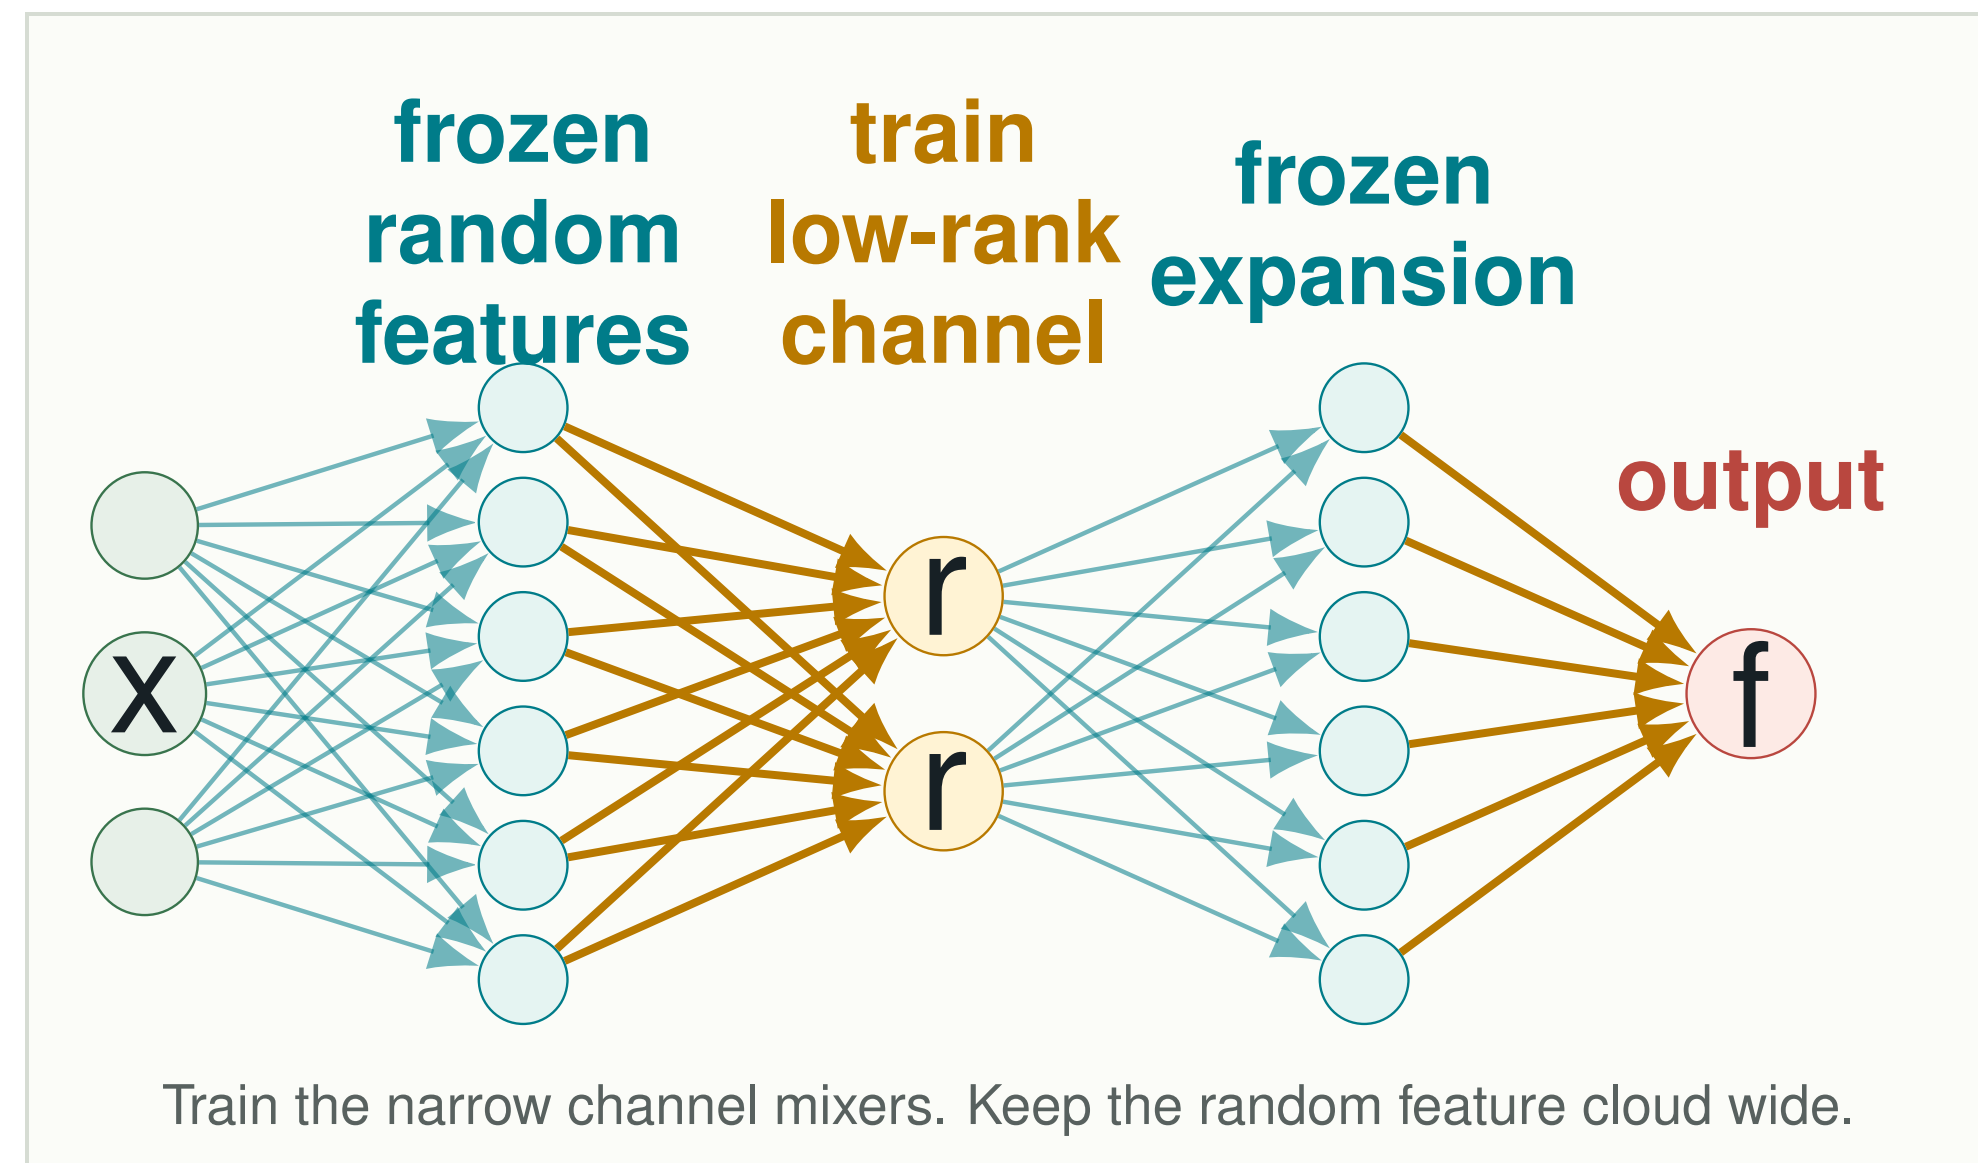
Low-Rank Networks Recover Weight & Functional Symmetry Better

ICML 2026 Workshop on Weight-Space Symmetries

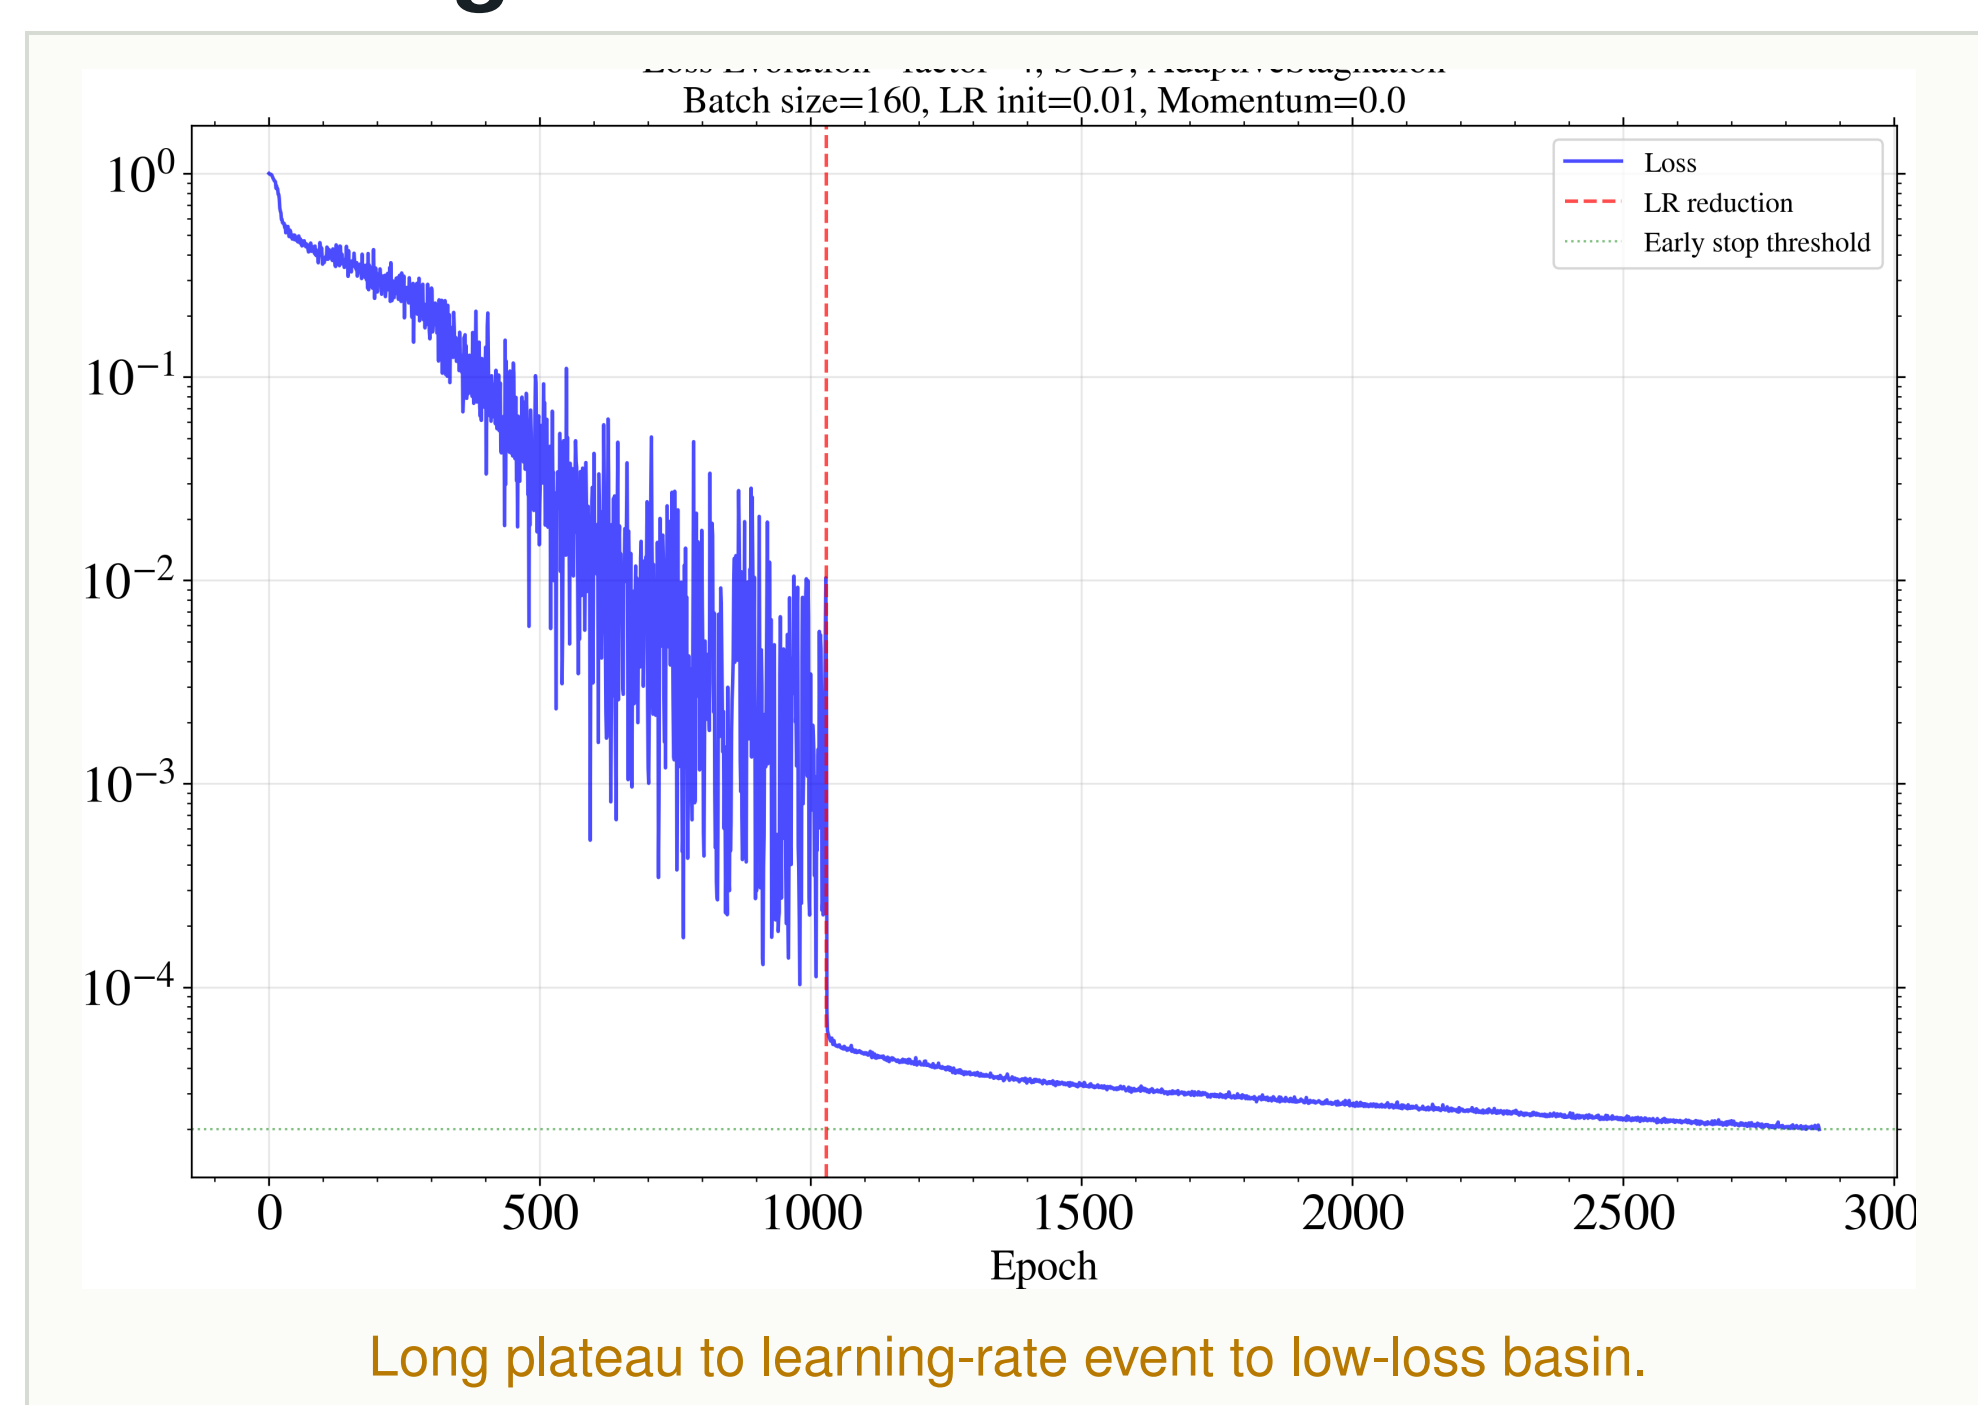
24 in x 36 in portrait

One visual message: Low rank makes training recover symmetry inside the network, not only at the output.

1. The Architecture



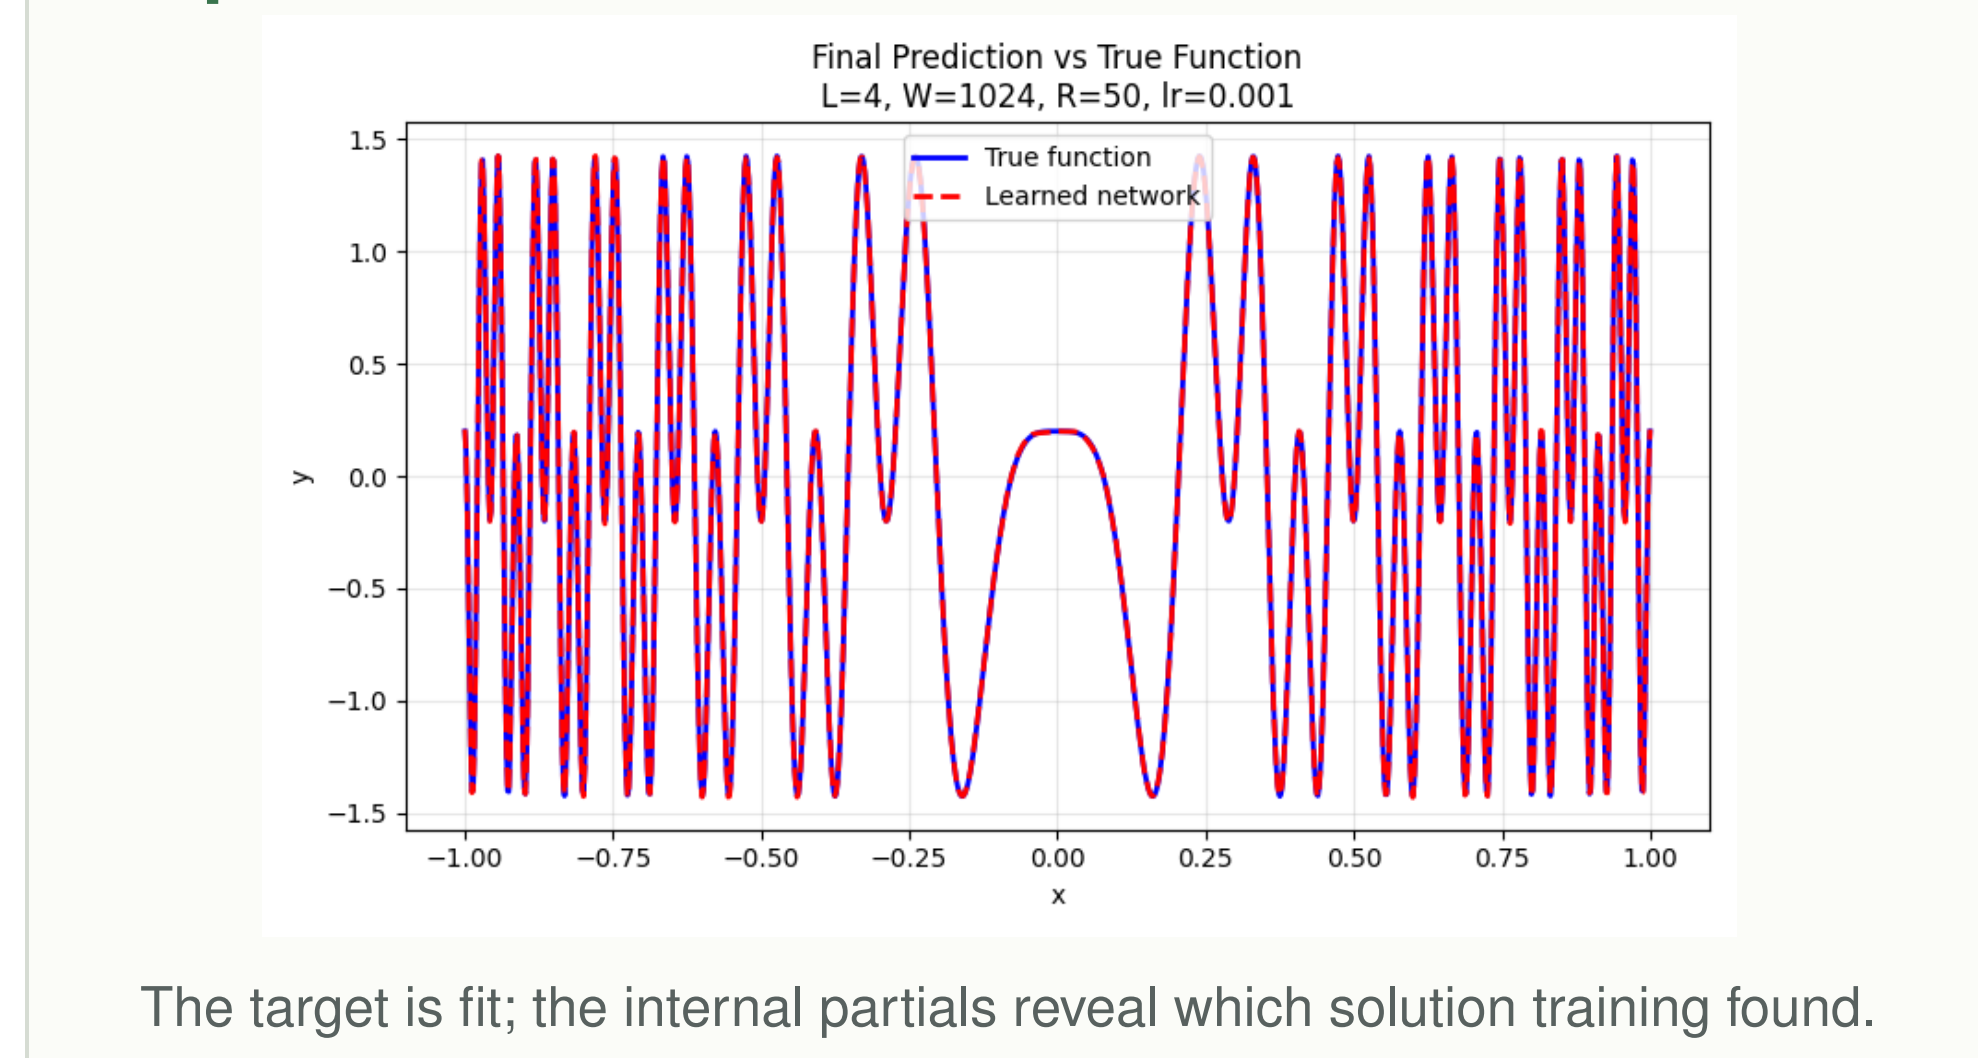
2. Training Finds A Plateau



What the optimizer seems to prefer


Symmetric low-rank partial minima appear to occupy a larger trainable volume than asymmetric hidden-feature minima.

Output fit alone hides the mechanism




3. The Partial Functions Become Symmetric

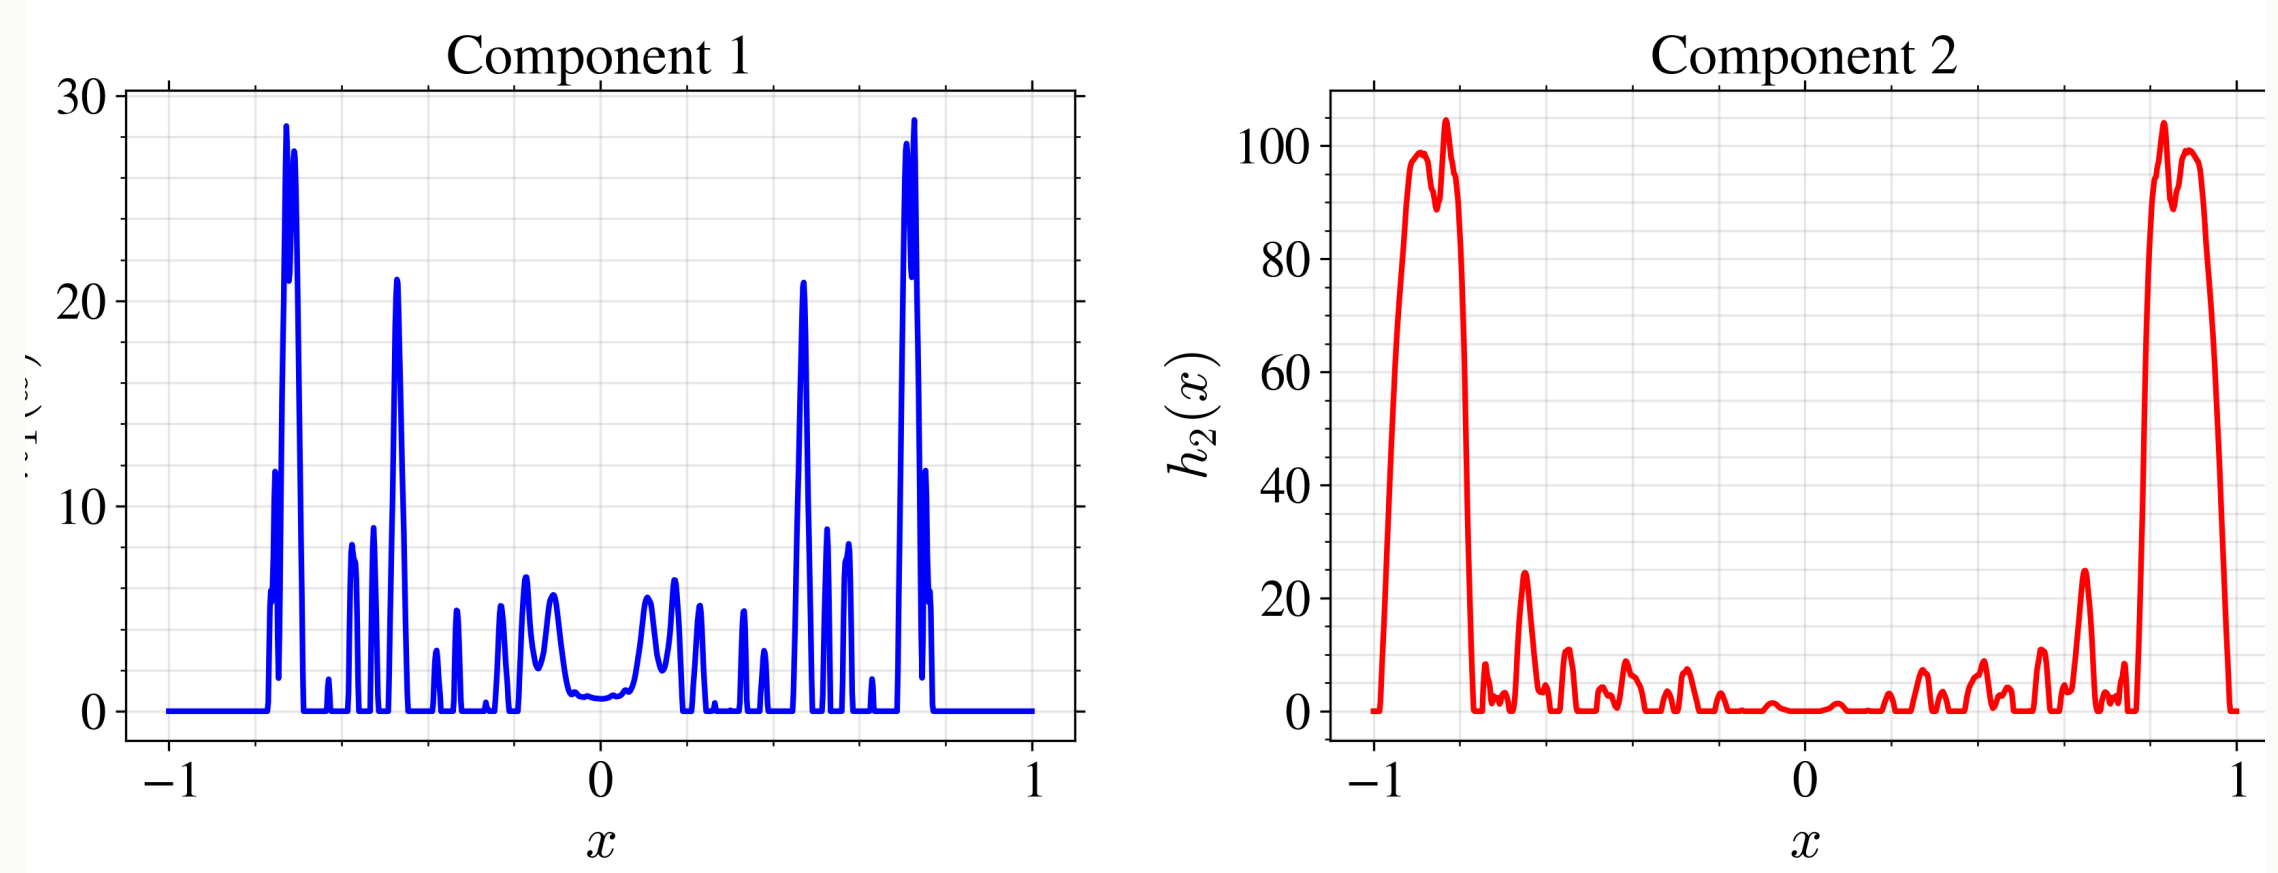
Early layer: smooth low-rank coordinates



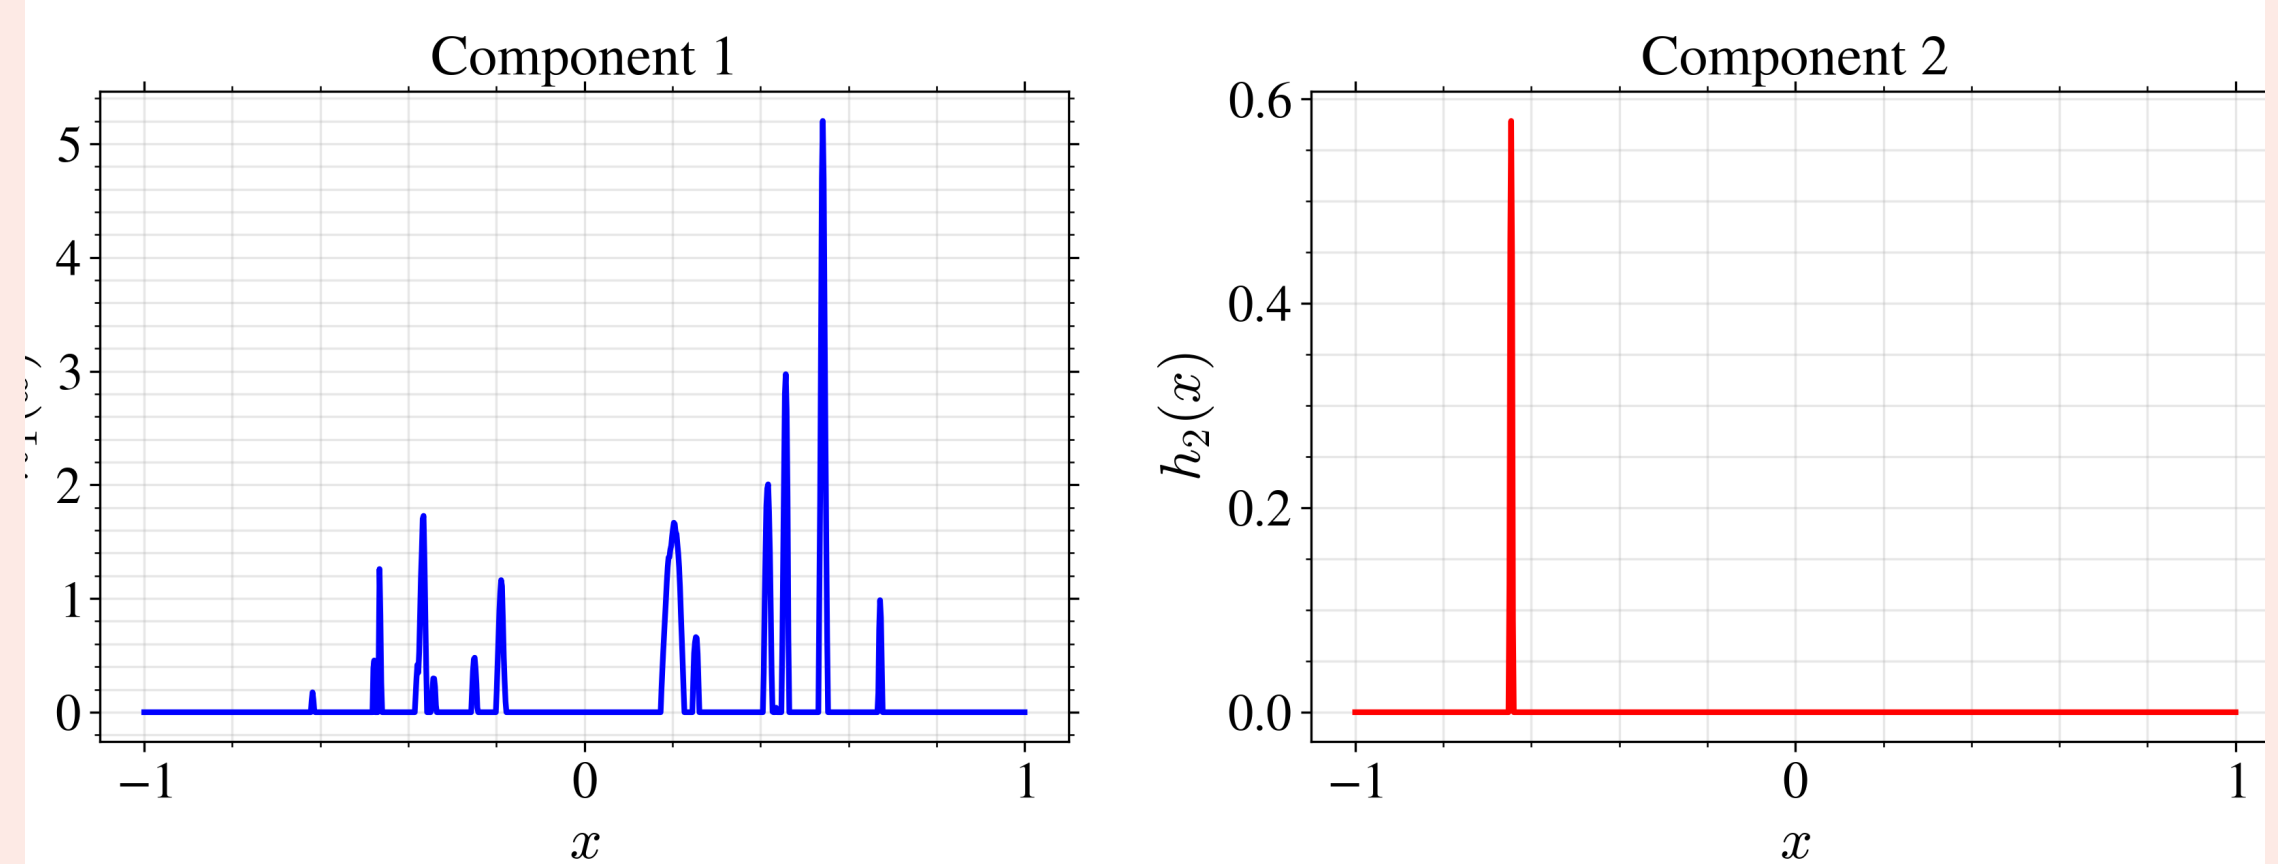
Middle layer: the mirror begins to lock in



Deep layer: sharp mirror structure



Dense MLP control



Same kind of target, hidden features stay asymmetric.

149

controlled reruns

$1e-5$

best test MSE

$\sim 1e-4$

active partial defect

output is not enough

Measure the inside

Active partials separate RF-LR from dense hidden features.

Why 1D?

After one dense layer, FCNN hidden states are already high-dimensional; 1D keeps the symmetry test clear.

Next theorem

Prove symmetric basins are wider for SGD/Adam.