AdaptiveMixGNN

Local Adaptive Inductive Bias for Heterophilic Node Classification

Miguel Alcocer Pérez, Javier Muñoz de Torres & Álvaro Morán Lorente · Universidad Rey Juan Carlos, Madrid, Spain

GRaM Workshop @ ICLR 2026 · Tiny Paper · Theme: Scale & Simplicity

github.com/miguelalcocker/AdaptiveMixGNN

The Problem

Most GNNs apply a uniform global filter to every node, assuming one dominant structural regime — but real graphs have regions that need fundamentally different treatments.

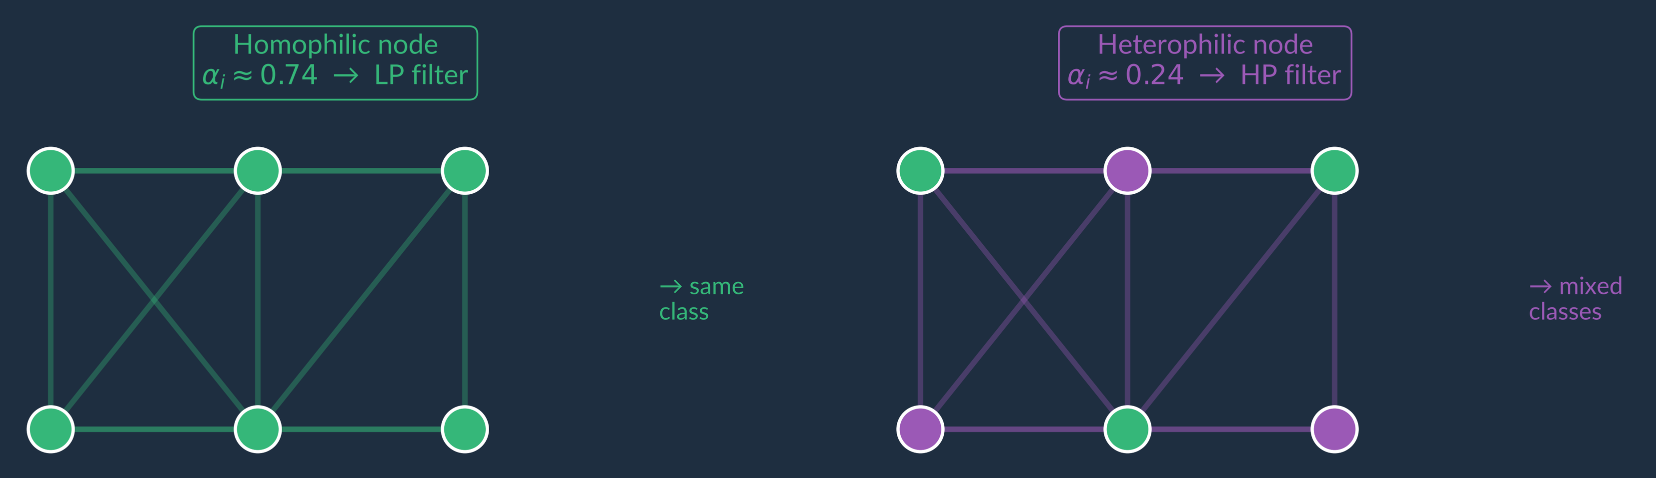

LP-only → Texas 57.8% | HP-only → Cora 24.2% | Optimal filter unknown a priori

The Solution

Each node learns its own mixing coefficient between S_{LP} and S_{HP}:

$$S_\alpha = \mathrm{diag}(\alpha)\, S_{\mathrm{LP}} + (\mathbf{I} - \mathrm{diag}(\alpha))\, S_{\mathrm{HP}}$$

$$\alpha_i = \sigma(\mathbf{h}_i^\top \theta + b)$$

$d+1$ parameters per layer
Minimal cost of local adaptivity

$O(|\mathcal{E}|)$ complexity
Identical to GCN — zero overhead

Permutation equivariant
Grounded in geometric deep learning

Time complexity (relative):

Each node learns its own spectral bias

Learned α_i correlates with local label structure — model discovers meaningful per-node frequency responses

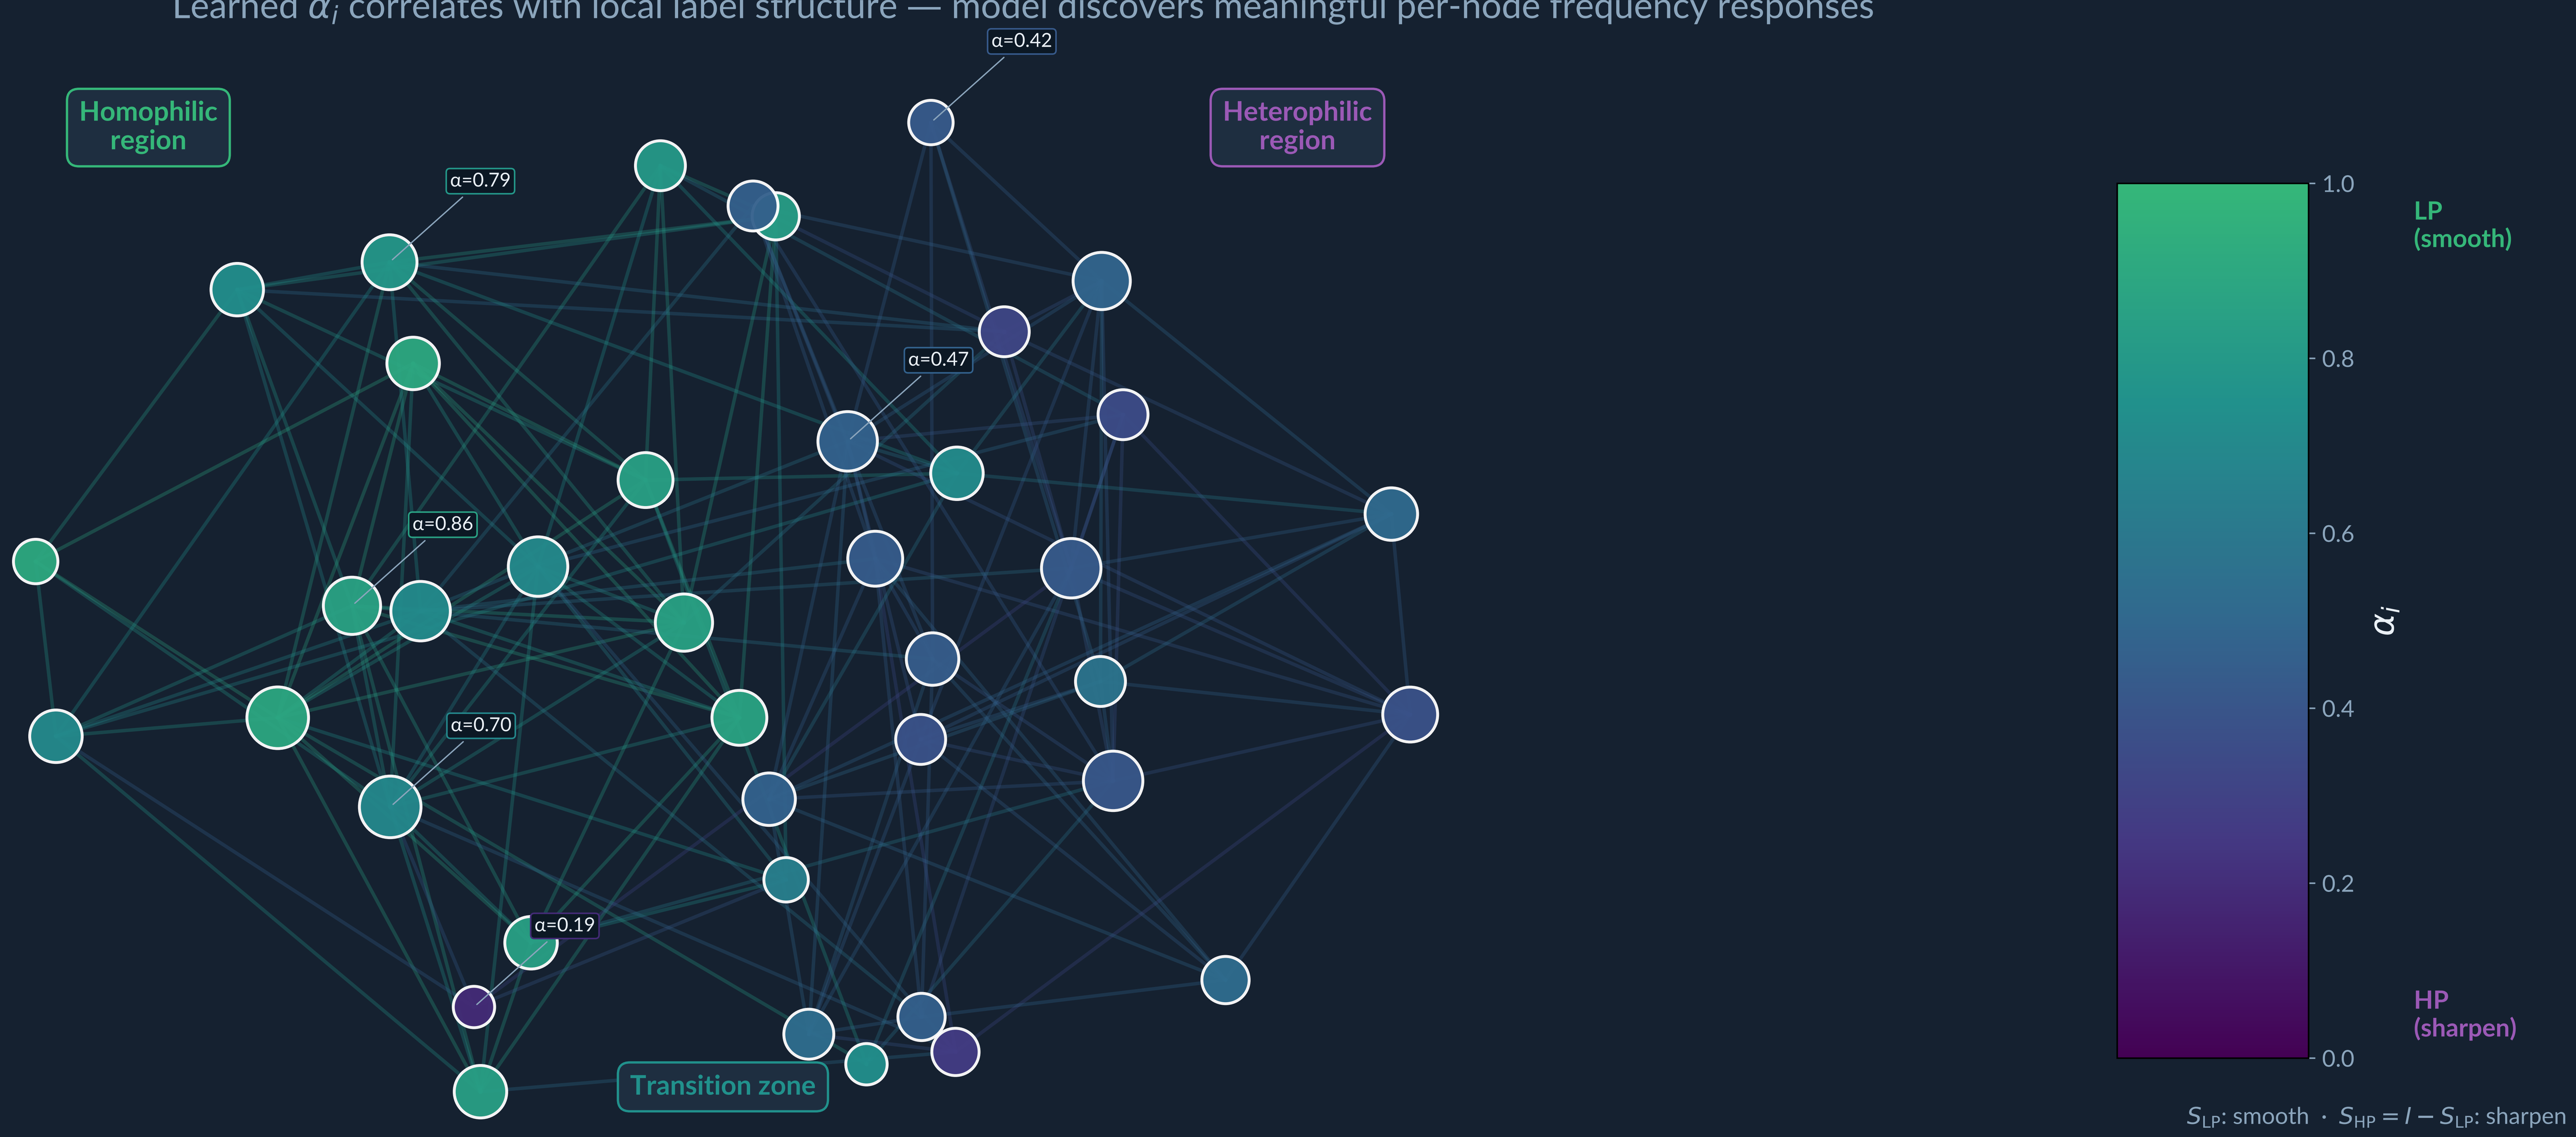

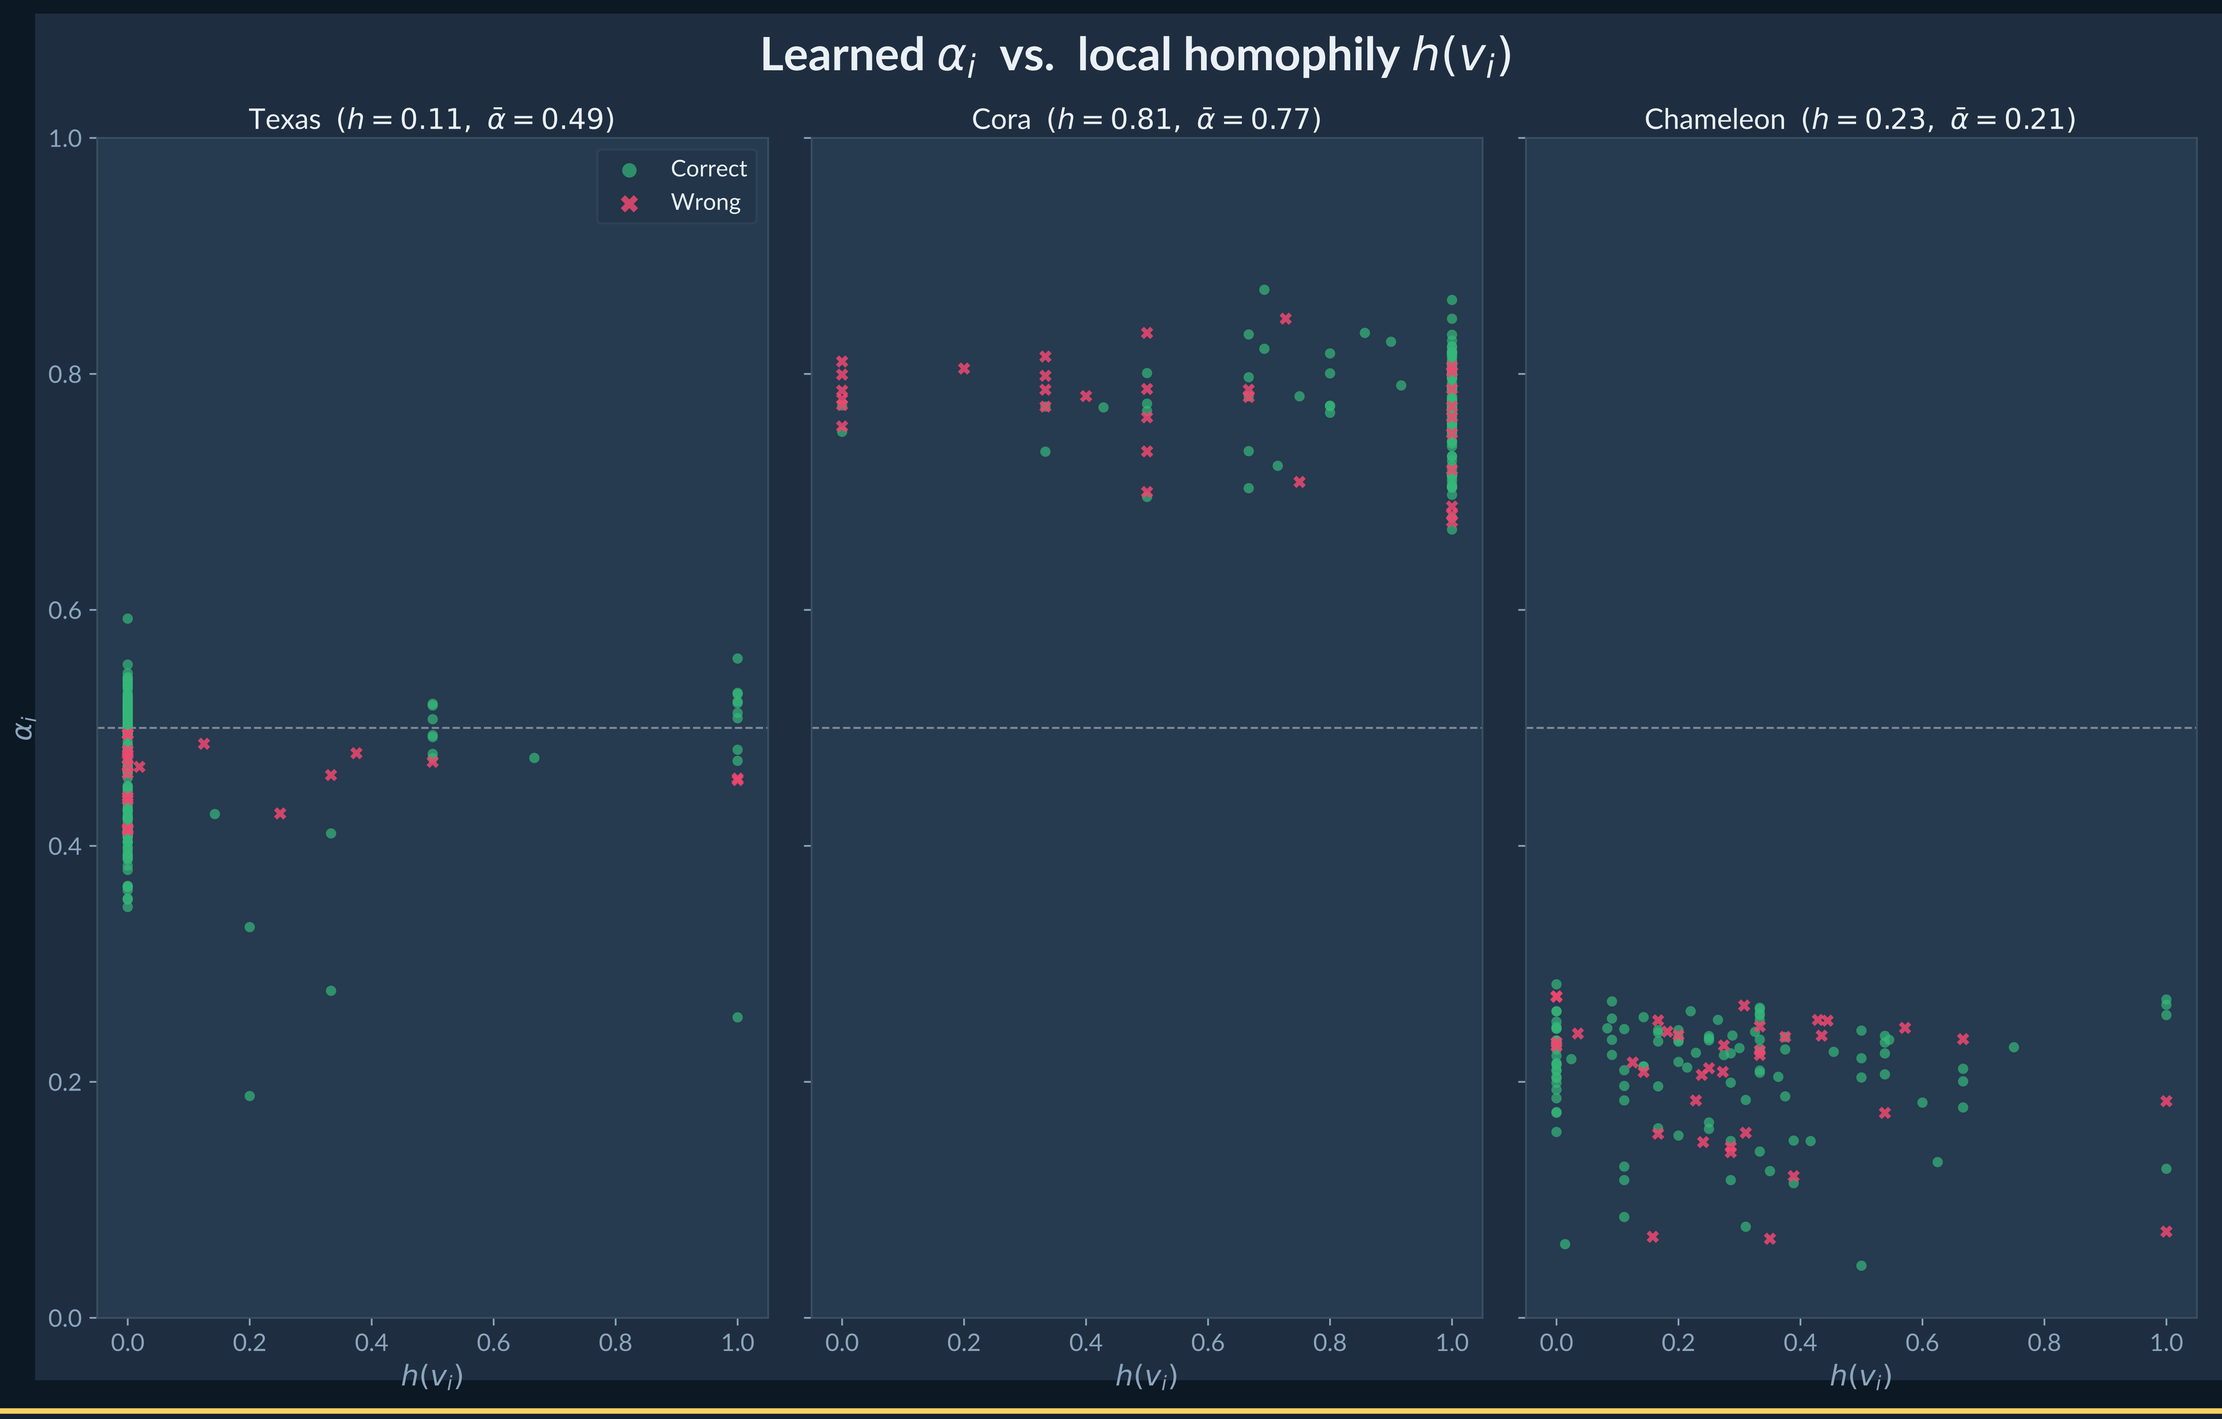

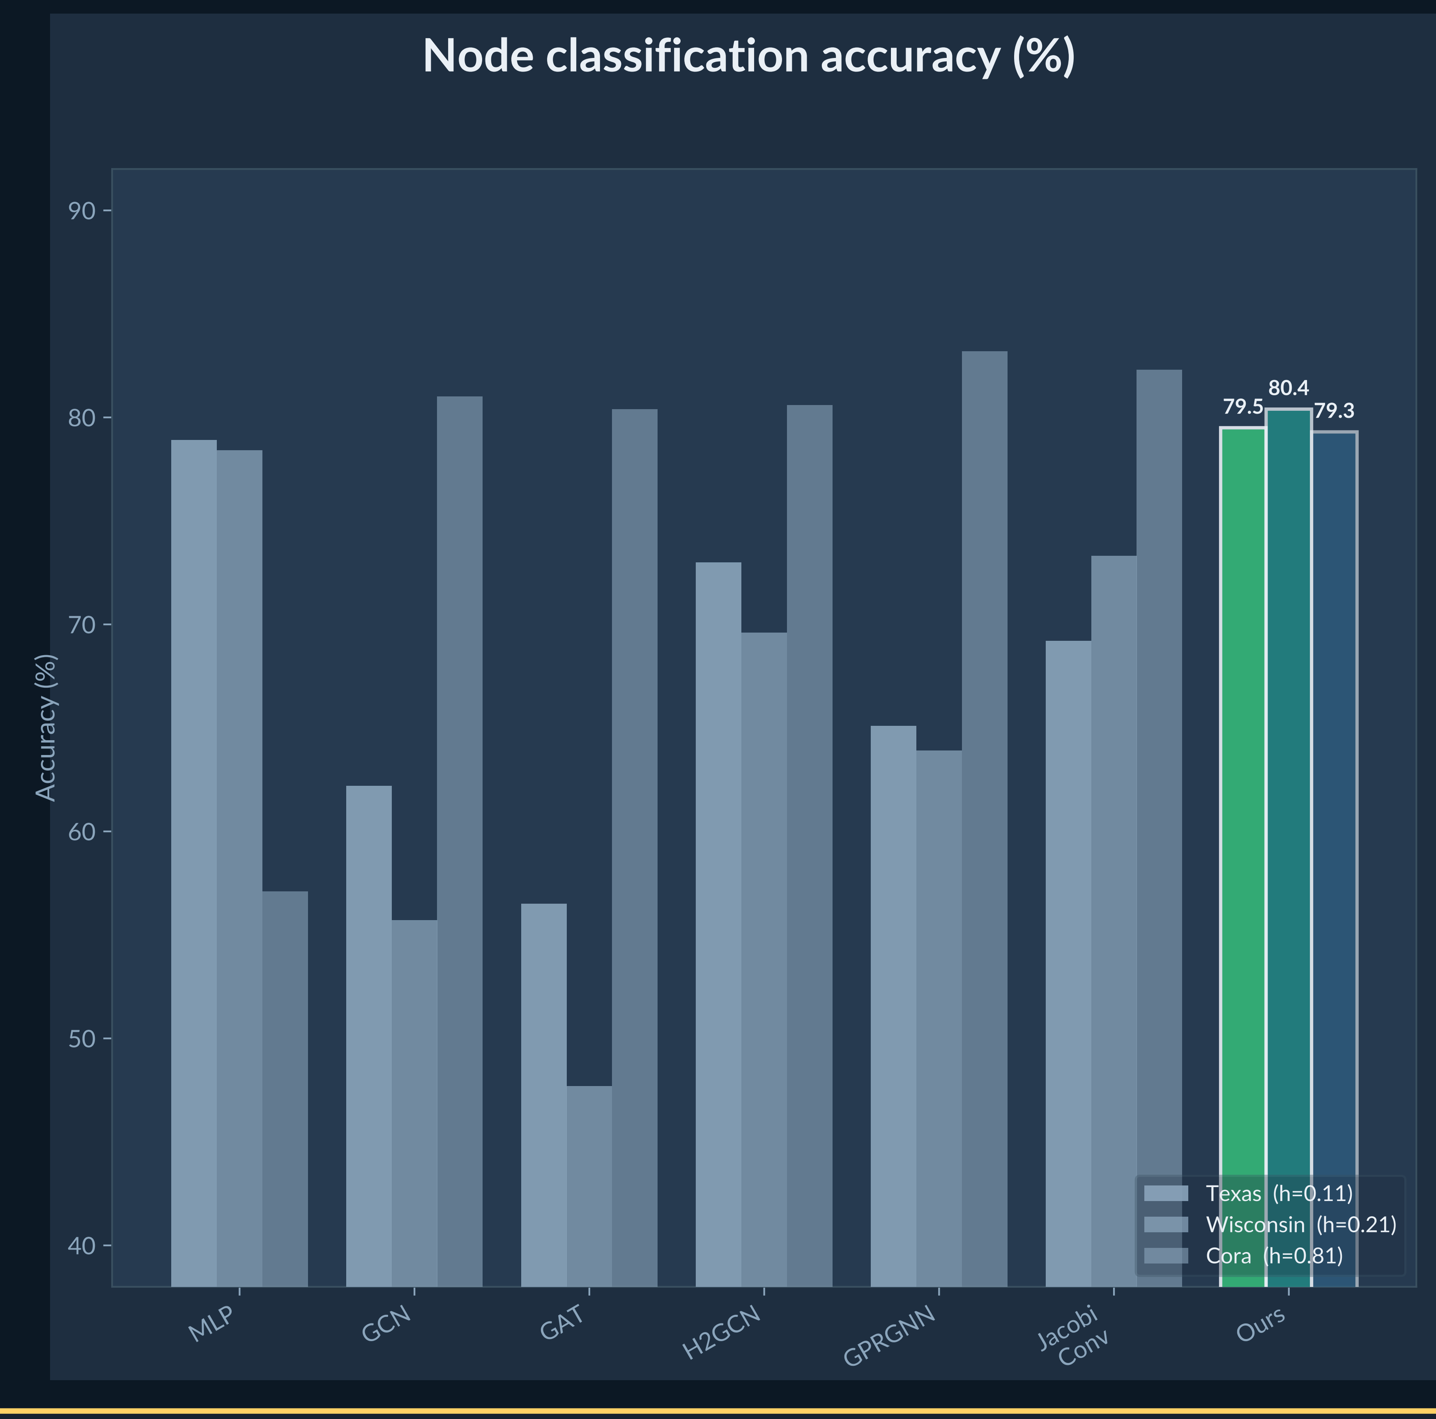

State-of-the-art among O(|E|) methods

79.5% Texas · 80.4% Wisconsin

GCN-like complexity with local adaptivity

d+1 params per layer · linear in edges

Grounded in geometry & spectral theory

Permutation equivariant · LP/HP filter bank

[1] Han et al. Node-MoE ICLR 2025 [2] Chien et al. GPRGNN ICLR 2021 [3] Wang & Zhang JacobiConv ICML 2022 [4] Platonov et al. ICLR 2023

Scan for code & data