Structure-Induced Information for Rerooting Levin Tree Search

Jake Tuero[1,2], Michael Buro[1], Laurent Orseau[3], Levi H. S. Lelis[1,2]

[1]Department of Computing Science, University of Alberta
[2]Alberta Machine Intelligence Institute (Amii), Edmonton
[3]Google DeepMind, London

May 30, 2026

Motivation

Hard planning problems need structure.

- Learned policies and heuristics help, but can struggle to reach distant goals.

Explicit subgoals are powerful but costly.

- They require generating or reconstructing intermediate target states.

Our idea:

Use **rerooting weights** to allocate search effort to useful internal roots, creating **soft subtasks** during search.

Levin Tree Search (LTS)

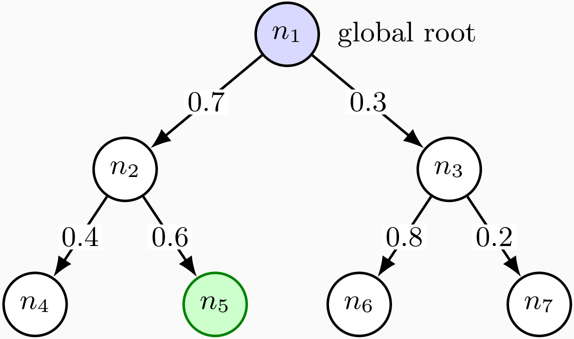

$$\varphi_{\mathrm{LTS}}(n) = \tfrac{d}{\pi}(n) = \frac{d(n)}{\pi(n)}$$

Search is guided only from the global root n_1.

$\sqrt{\text{LTS}}$ Intuition

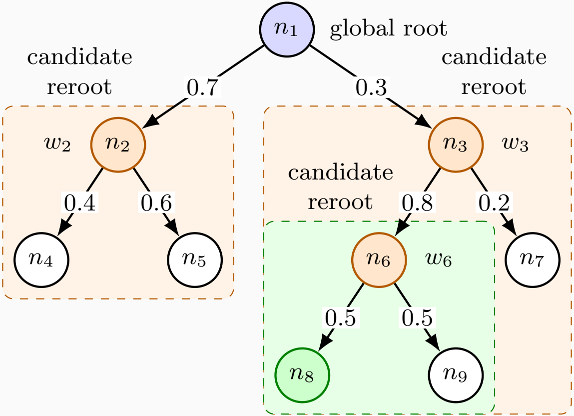

$$\varphi_{\sqrt{\text{LTS}}}(n) = \min_{n_t \prec n} \frac{1}{w_t} \, \frac{d}{\pi}(n_t \prec n)$$

Search effort is redistributed across multiple local roots.

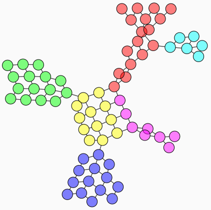

Set color counts:

1. Select graph $G_k = (V_k, E_k)$ for a given cluster level $k \in [1, N]$.
2. Color nodes $v \in V_0$ using cluster association from V_k:

M_{τ,c_i} = {
[cyan] = 6, [red] = 17, [yellow] = 12,
[magenta] = 8, [blue] = 14, [green] = 15
}

$$w_t = \frac{1}{M_{\tau,c_t} + \delta_{\tau,c_t}}$$

Global Structure Rerooter

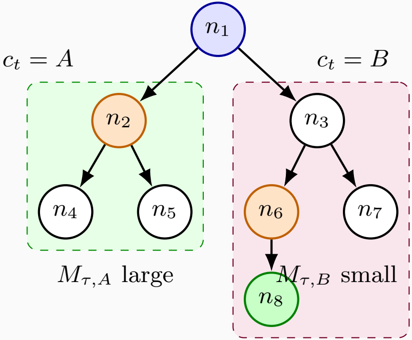

$$w_t^L = \frac{1}{M_{\tau,c_t} + \delta_{\tau,c_t}}$$

Prefer less-explored structural regions.

$\sqrt{\text{LTS}}$-H Heuristic-Based Rerooter

Local Heuristic Rerooter

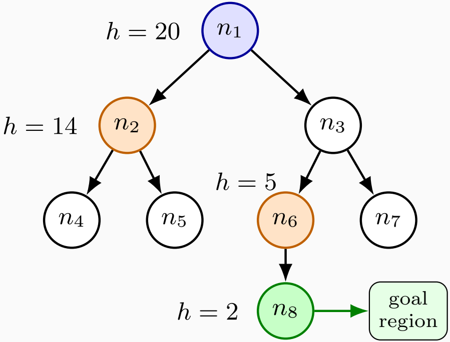

$$w_t^H = \exp\left(-\alpha \frac{h(n_t)}{h(n_1)}\right)$$

Prefer roots that appear closer to the goal.

$\sqrt{\text{LTS}}$-LH Hybrid Additive Rerooter

Hybrid Additive Rerooter

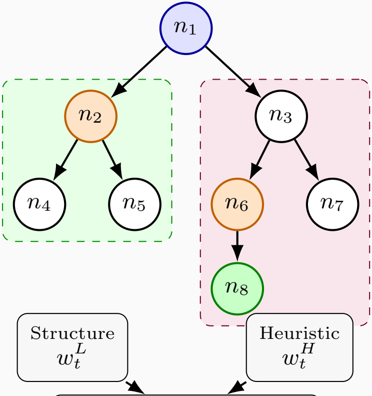

$$w_t^{LH} = u_L w_t^L + u_H w_t^H$$

Use global structure for broad allocation and heuristic progress for local refinement.

Results - Training Efficiency

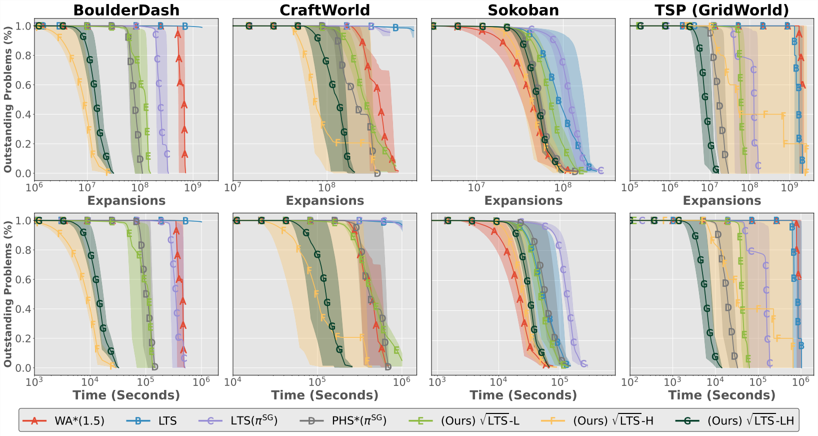

Results - Domain Complexity Scaling

Algorithm	Expansions	Time	Solved	Nodes/s
	BoulderDash (10%)			
PHS*(π^{SG})	30,686,786	10.63	10,000	802.1
$\sqrt{\text{LTS}}$-LH	18,654,789	5.36	10,000	967.4
	BoulderDash (20%)			
PHS*(π^{SG})	299,672,516	137.12	10,000	607.1
$\sqrt{\text{LTS}}$-LH	25,971,412	7.70	9,992	936.6
	BoulderDash (30%)			
PHS*(π^{SG})	429,407,720	278.23	11	428.7
$\sqrt{\text{LTS}}$-LH	57,780,467	16.50	10,000	972.5
	BoulderDash (40%)			
PHS*(π^{SG})	—	—	—	—
$\sqrt{\text{LTS}}$-LH	75,672,897	22.49	9,995	934.79
	BoulderDash (50%)			
PHS*(π^{SG})	—	—	—	—
$\sqrt{\text{LTS}}$-LH	97,117,682	30.13	9,990	895.2

Takeaway

Rerooting provides a scalable way to exploit structure in policy-guided tree search without explicitly generating subgoals.

Structure from search itself.

- Use clusters, heuristics, or both to define rerooting weights.

Subgoal-like behavior without subgoal models.

- Allocate effort to useful internal roots instead of reconstructing target states.

Practical and scalable

- Improves online training efficiency across challenging planning domains.