DEPTH: DEL Epistemic Planner with Tier Heuristics

Halim Djerroud
Université Paris-Saclay, UVSQ, LISV, 78140, Vélizy-Villacoublay, France
halim.djerroud@uvsq.fr

IPC 2026 – Epistemic Planning Track

Abstract

DEPTH (DEL Epistemic Planner with Tier Heuristics) is our entry to the first Epistemic Planning Track of the International Planning Competition 2026. It is a DEL-based planner built on top of the `plank` toolkit, supporting EPDDL specifications across the basic, intermediate, and hard levels.

The core algorithm, Epistemic H*, is a two-queue best-first search adapted to Dynamic Epistemic Logic. A primary queue (OPEN) advances toward the goal by prioritising successors that strictly improve the heuristic; a secondary queue (INCONS) retains the rest. The role of the secondary queue is to preserve *setup actions*—announcements, sensing steps, signal-passing actions—that fail to reduce the heuristic locally but enable later goal-achieving actions. When OPEN becomes empty, INCONS is flushed back into it, ensuring that no reachable state is discarded.

Goal progress is estimated by a deliberately coarse heuristic, the *tier*, which counts the number of unsatisfied subformulas of the goal after decomposition. Multi-agent modal goals are split into per-agent singletons, yielding an over-approximation that is sound for the heuristic but not for the goal test, which always uses the original formula. The tier is inadmissible by design, trading optimality for informativeness; the two-queue mechanism is what underwrites completeness despite this. Bisimulation contraction is applied to every generated successor before insertion into the frontier, so that logically indistinguishable states are identified regardless of how they were reached.

We prove that Epistemic H* is complete for finite DEL planning tasks. Preliminary experiments on the IPC 2026 sample benchmarks show that DEPTH significantly outperforms the reference BFS planner on instances with large Kripke models, solving Grapevine (32 worlds) in milliseconds where BFS exhausts the 60-second timeout. We discuss the canonical inadmissibility cost (one-step longer plans on selected Coin-in-the-Box instances) and outline directions for stronger heuristics in the extended post-competition version of this work.