





Mind the Gap: Promises and Pitfalls of Hierarchical Planning in LeWorldModel



Niccolò Caselli* Francesco Massafra* Samuele Punzo*
Salvatore Lo Sardo* Ippokratis Pantelidis* Sathya Kamesh Bhethanabhotla†
University of Amsterdam *Equal contribution † Supervisor


Why Hierarchy? Long-Horizon Planning Is Hard

- **Motivation:** LeWM plans directly over primitive-action sequences. Farther goals make CEM search harder and compound model error.
- **Our contribution:** combine LeWM with Zhang et al.'s latent hierarchy as a frozen low-level retrofit, then diagnose when it fails and when it works.

Frozen low-level LeWM + latent hierarchy


LeWorldModel + Hierarchical Planning ⇒ Hi-LeWM
 Maes et al. (2026) Zhang et al. (2026) Caselli et al. (2026)

Background: LeWorldModel



- **Flat LeWM:** CEM/MPC searches primitive action sequences with latent rollouts.
- **Long horizons:** farther goals increase optimizer burden and model-error compounding.
- **Hi-LeWM idea:** choose latent waypoints first, then let frozen low-level LeWM execute toward them.

Method: Add hierarchy, freeze LeWM



- **Freeze** the LeWM encoder, low-level predictor, and low-level action encoder.
- **Train** a macro-action encoder g and high-level predictor p_{hi} over latent waypoints.
- **Plan** high-level subgoals; let frozen low-level LeWM execute primitive actions toward them.

Hierarchy Is Not Plug-and-Play

PushT offset	Flat LeWM	Naive Hi-LeWM	Gap
$d = 25$	94.0%	89.3%	-4.7 pp
$d = 50$	52.7%	38.7%	-14.0 pp
$d = 75$	18.0%	15.3%	-2.7 pp

Takeaway


Naive hierarchy underperforms flat LeWM.

Where does the Hierarchy Break?

Oracles

Oracle acting reaches $73.3 \pm 4.2\%$ at $d = 50$: frozen LeWM can execute good intermediate targets.

Subgoals



Oracle targets are coherent; CEM targets are diffuse and misaligned with reachable states.

CEM



Offline, CEM improves the terminal objective while breaking waypoint 1.

Latent MSE ↓	Teacher-forced	Open-loop true	Open-loop CEM
Stage 1	0.081	0.081	0.347
Stage 2	0.104	0.216	0.011

Takeaway

CEM leaves the data-supported macro-action space.

Why this happens



Intervention: Put CEM on Rails


Empirical-macro CEM

Empirical-macro CEM narrows the optimizer's playground: sample a macro-action from real trajectories, then optimize only a small residual:


$$\xi = \xi_{\text{bank}} + \lambda_{\text{res}} \epsilon.$$

This keeps high-level search near behavior the dataset actually supports.

Empirical-Macro CEM Produces Sharper Subgoals



Results



Takeaway

On PushT, constrained search gains at $d = 50$ and $d = 75$.

Final Takeaways

- **Not plug-and-play:** naive hierarchy underperforms flat LeWM.
- **Failure mode:** CEM leaves the data-supported macro-action space.
- **PushT result:** constrained search recovers gains at $d = 50$ and $d = 75$.

References. Maes et al. (2026), *LeWorldModel*; Zhang et al. (2026), *Hierarchical Planning with Latent World Models*; Rubinstein (1999), *Cross-Entropy Method*.