

Learning When to Reason


Same calls. Better timing. More events covered.

Question. With a fixed number of reasoning requests, which moments should receive them?


Approach. Learn pre-event opportunity scores; compare constrained call sets at exact call parity.

Predict opportunities; maintain control

Selective scheduling pipeline



Illustrative Town10HD timing replay



Recent state history and compact BEV context inform opportunity scores. The fast policy continues controlling the agent; reactive recovery follows an independently accounted path.


The timing images are illustrative frames from a separate visualization run.

What is measured?

Useful-Window Coverage (UWC) is the fraction of distinct event-specific pre-onset windows reached by at least one call.

Frozen trajectories and identical call counts isolate placement quality. Backend answers, online commitment, and closed-loop utility are separate questions.

165 calls: 23 event windows versus approximately six



+26.71 percentage points

UWC@0.5 s at the 2% budget

95% bootstrap interval: 12.64 to 40.97 percentage points. The primary Town10HD comparison uses a development-selected checkpoint; Random-k is the strongest non-oracle baseline at this budget.

On fresh Town10HD scenarios, a separately frozen development-selection procedure yields +0.2392 normalized log-budget UWC-AUC over Random-k, with positive differences at all five budgets.

Selected-model evidence does not establish arbitrary-seed stability. Complete seed distributions and matched-search comparisons are retained in the paper and data.